AD-A168 613 VISUAL REPRESENTATIONS SUBSERVING TEXTURE PERCEPTION (U) OREGON UNIV EUGENE J BECK ET AL. 15 JAN 86 AFOSR-TR-86-0290 F49620-83-C-0093 1/2

UNCLASSIFIED F/G 5/10 NL

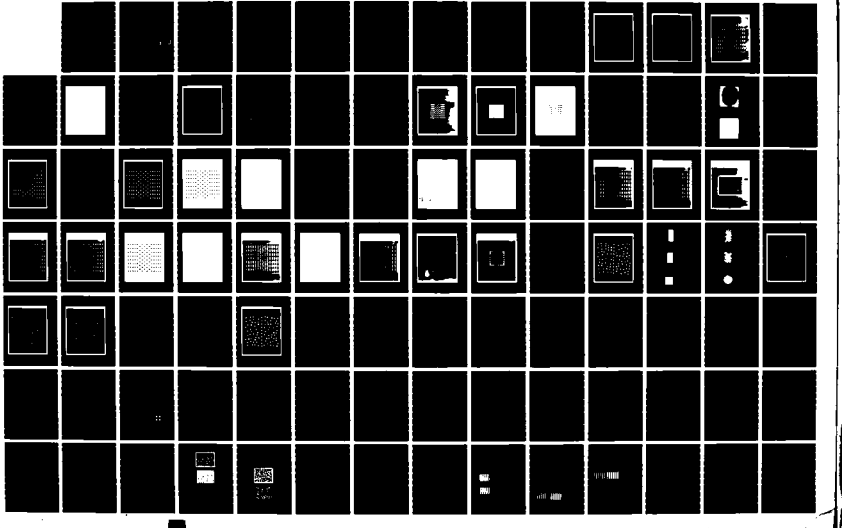

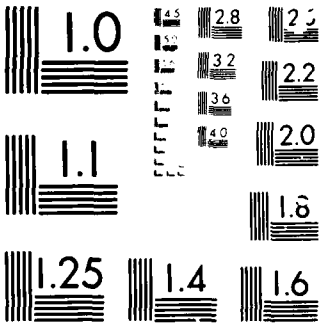

VISUAL REPRESENTATIONS SUESERVING TEXTURE PERCEPTION

Final Report

by

Jacob Beck

Department of Psychology

Kent A. Stevens

Department of Computer and Information Science

University of Oregon

Eugene, Oregon 97403

January 15, 1986

Prepared for:

Air Force Office of Scientific Research
Building 410
Bolling AFB, DC 20332-6448

Contract Number: F1962083-C-0093

AFOSR Project Manager: Dr. John Tangney

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AFOSR-TR- 86-0290	2. GOVT ACCESSION NO	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Visual representations subserving texture perception		5. TYPE OF REPORT & PERIOD COVERED Final Report April 1, 1983-August 31, 1985 6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Jacob Beck, Kent A. Stevens		8. CONTRACT OR GRANT NUMBER(s) F49620-83-C0093
9. PERFORMING ORGANIZATION NAME AND ADDRESS University of Oregon Eugene, Oregon 97403		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61102F 2313/F5
11. CONTROLLING OFFICE NAME AND ADDRESS USAF, AFC AFOSR Building 410 Bolling AFB DC 20332		12. REPORT DATE January 15, 1986 13. NUMBER OF PAGES 105
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Vision, texture perception, texture segmentation

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

We have conducted research on the role of spatial filtering, features, and grouping in texture segregation. Our experiments indicate the interplay of two different processes. One process involves the differential excitation of elongated receptive fields. Texture segregation is a function of energy differences (contrast and size) that are largely extracted from the lower spatial frequencies. The second process involves local processes of linking between localized features. Linking of contours, for example, is a function

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

of contour smoothness, collinearity, orientation, etc. These effects cannot be explained in terms of low frequency differences. Studies of the linking of discrete textures have provided convergent evidence for explicit place markers and the role of similarity of attributes such as color and contrast in establishing these groupings. We have also examined the role of pairwise linkings, or virtual lines for imposing global organization on the localized intensity changes. Also, at the level of contour representation within texture, we have shown the role of the concave cusp, a localized geometric feature, in determining figure-ground assignment in texture.

Contents

ABSTRACT

We have conducted research on the role of spatial filtering, features, and grouping in texture segregation. Our experiments indicate the interplay of two different processes. One process involves the differential excitation of elongated receptive fields. Texture segregation is a function of energy differences (contrast and size) that are largely extracted from the lower spatial frequencies. The second process involves local processes of linking between localized features. Linking of contours, for example, is a function of contour smoothness, collinearity, orientation, etc. These effects cannot be explained in terms of low frequency differences. Studies of the linking of discrete textures have provided convergent evidence for explicit place markers and the role of similarity of attributes such as color and contrast in establishing these groupings. We have also examined the role of pairwise linkings, or virtual lines for imposing global organization on the localized intensity changes. Also, at the level of contour representation within texture, we have shown the role of the concave cusp, a localized geometric feature, in determining figure-ground assignment in texture.

1. INTRODUCTION

Our overall goal is to develop a model for the representation and processing of texture in human vision. In examining some fundamental perceptual tasks including the detection of texture differences and the detection of geometric structure in texture we have found convergent evidence for two basic types of processing, which implicate two very different types of texture description. On the one hand some tasks appear to be performed in the energy domain, where variables such as contrast and size dominate. Generally for these tasks a representation of intensity changes in terms of low spatial frequencies seems important. On the other hand, we have identified geometric tasks for which the localization of intensity changes and similarity of their attributes is important. These tasks require the explicit representation of position and attribute, and in contrast to the former, appear to be independent of low spatial frequency content. The subsequent sections of this final report describe in detail the specific energy and geometric characterist that we have found salient. The organization of the report is sketched in the following.

Section 2 reports experiments that use a competition paradigm to investigate texture segregation. The experiments present evidence for a two component theory of texture segregation. The first component is an operator sensitive to the outputs of spatial frequency sensitive mechanisms. This operator segregates regions acording to differences in contrast sign and differences in low spatial frequencies. There is no interaction between positive and negative contrasts, and strong segregation can occur when there is only a small intensity difference between texture elements of opposite contrast. If the contrast sign is the same, texture segregation occurs only if the shapes of the distributions of low spatial frequency mechanisms responding differ. Thus, equal size texture elements differing in contrast magnitude fail to cause texture segregation unless the contrast difference is very large. Low frequency mechanisms may also mask high frequency mechanisms. A small size and a high contrast that stimulate the same low frequency mechanisms as a large size and a weak contrast fail to cause texture segregation. Hue is a weak feature relative to contrast. That is, hue differences are not sufficient to cause texture segregation in the presence of contradictory information from contrast mechanisms. The second component is an operator that agregates texture elements on the basis of geometric descriptors such as aspect ratio, alignment, contour smoothness etc., and then segregates regions on the basis of "emergent features". The effects of these variables cannot be explained in terms of diferences in the response of low spatial frequency mechanisms. The linear organization appears to be derived from a nonlinear operator that connects nearby localized spatial tokens.

Section 3 provides further evidence for explicit grouping tokens. The evidence is provided, in part, by experiments that show groupings in dot patterns that are devoid of low spatial frequency content. Additionally, a rivalry paradigm is used to demonstrate the importance of similarity of various attributes, giving experimental results that are consistent with the notion of explicit grouping tokens, and contrary to the predictions of various alternative schemes. A

computational model of the linking process is examined and demonstrated on various dot patterns.

Section 4 reports our finding that figure may be distinguished from ground on the basis of information provided by concave cusps in image curves. Concave cusps are prevalent in natural textures, occurring in images wherever the silhouettes of convex, opaque objects abut or partially occlude one another. The points of contact between silhouettes present concave cusps, each indicating the local assignment of figure versus ground across the contour. There are known tendencies to interpret as figure those regions that are lighter, or smaller or more convex. We show that the concave cusp is another, distinct, determiner of apparent figure-ground that provides a stronger influence on figure-ground than lightness or size. While concave cusps are associated with overlapping convex figures, the salience of the cusp appears to derive from the local geometry and not from the adjacent contour convexity.

2. TEXTURE SEGREGATION

2.1 Introduction

Agreement exists that texture segregation derives from the preattentive perceptual processes that extract information about simple properties such as orientation, size, brightness, hue, line endings, etc (Beck, 1966, 1972, 1982; Julesz, 1978; Julesz & Schumer, 1981). Preattentive texture segregation, however, depends not only on the first-order statistics of these properties, but also on the orientation, length, and certain other attributes of organized structures(Beck,1983; Beck et al., 1983). Figure 1 shows three regions of five rows each. In the top and bottom regions the squares are arranged in vertical stripes. In the middle region the squares are arranged in diagonal stripes. The middle region is segregated from the top and bottom regions, a perception we call tripartite segregation. Two different processes can be involved in the detection of such structures. One is a process that detects the differential excitation of visual channels. By visual channels we mean the range of sizes of elliptical receptive fields. These do not differentiate texture elements from each other or from the background. They respond to total stimulus energy resulting from averaging a stimulus over different size spatial regions. The primitives for tripartite segregation are the outputs of the visual channels. The second is a process that detects differences in the geometric properties of texture elements. Examples of such properties are contour smoothness, contour misalignment, orientation, size, etc. The effects that these properties have can not be explained in terms of low frequency energy differences that result from averaging intensities over spatial regions.

We will first report studies which describe how energy variables such as intensity, size, spacing, and hue of the squares and the background affect tripartite segregation. Second, we will consider the extent to which the effects of energy variables can be explained in terms of the differential stimulation of elliptical receptive fields when a display is convolved with a difference of Gaussians. Third, we will present data from a different task, in which subjects are required to detect a line in a random field. This task puts more emphasis on the processes that make a larger unit (a line) out of smaller elements and less emphasis on noticing what elements are the same or different. Texture segregation occurs as a result of a process which explicitly knits together, via some nonlinear operation, the local responses to the geometric properties of the texture elements.

2.2 Tripartite Experiments 1-5: Equal Size Squares

Figure 2 shows that adding lighter squares can destroy tripartite segregation. One no longer sees that in the top and bottom regions the squares are arranged in verical stripes, while in the middle region, the squares are arranged in a checkerboard. There is an impression of overall uniformity. Figure 3 shows that when the background intensity is between instead of above the intensities of the squares, a stable tripartite segregation occurs. Experiments 1 through 5 investigated how the intensity, spacing, and hue of the squares and background affect tripartite segregation.

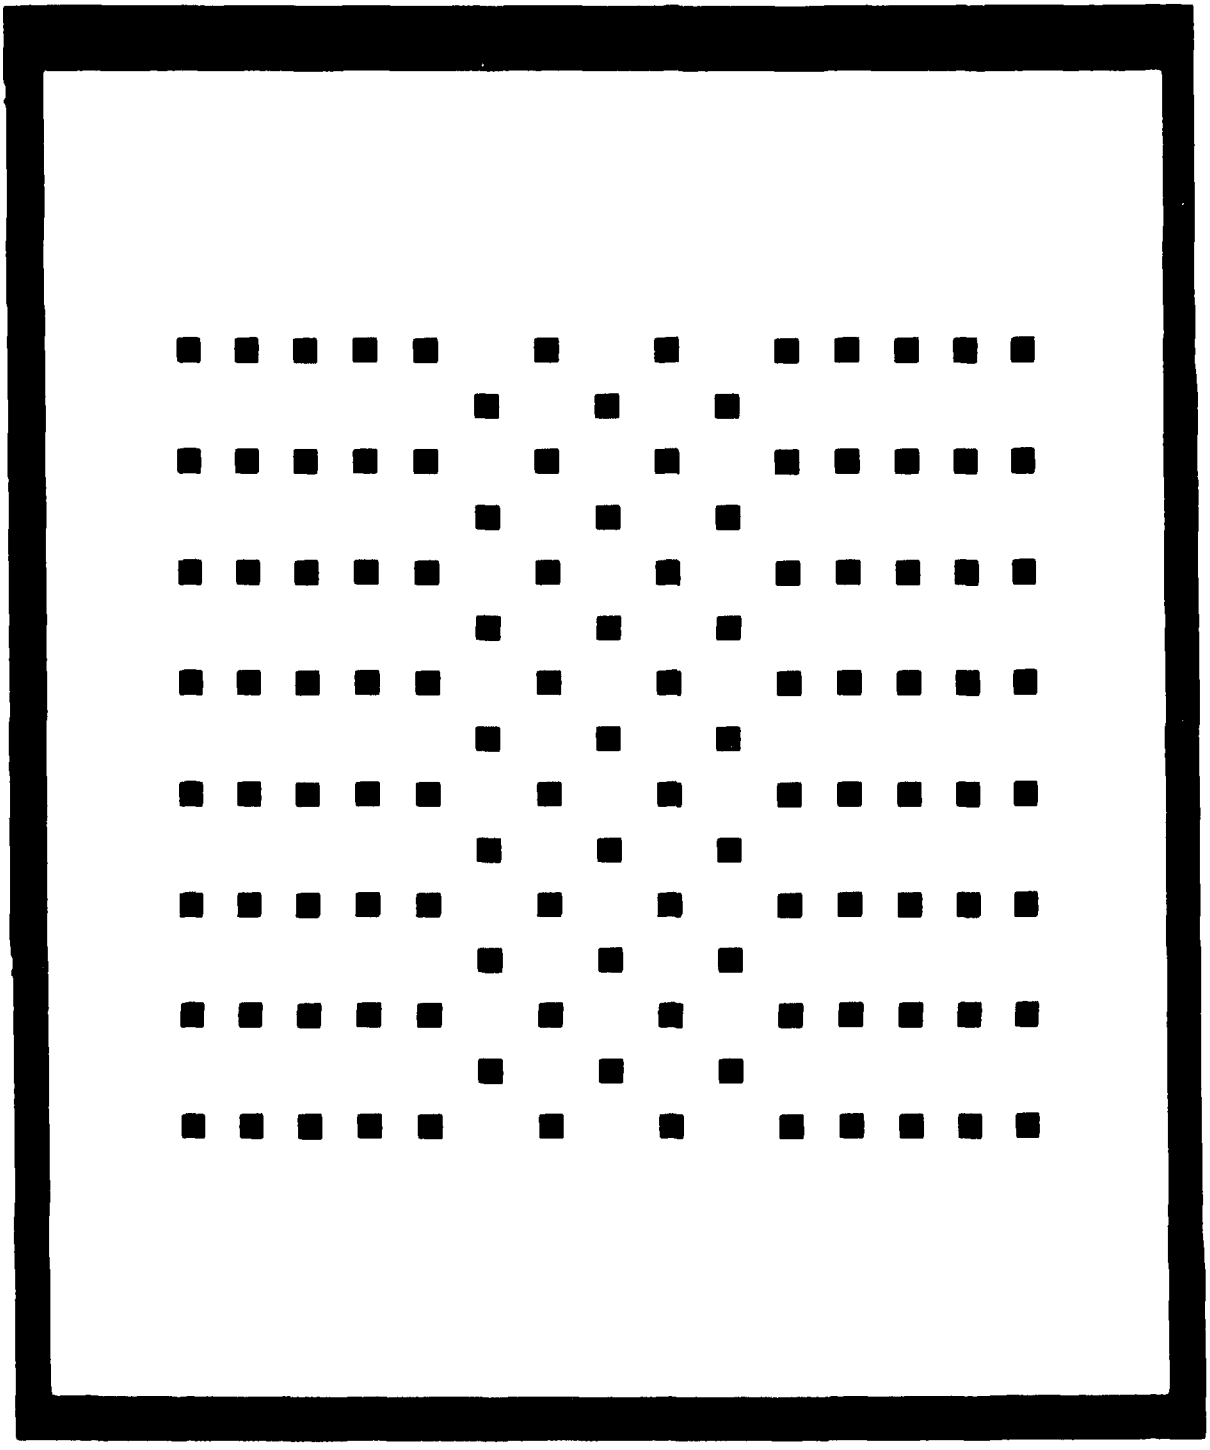

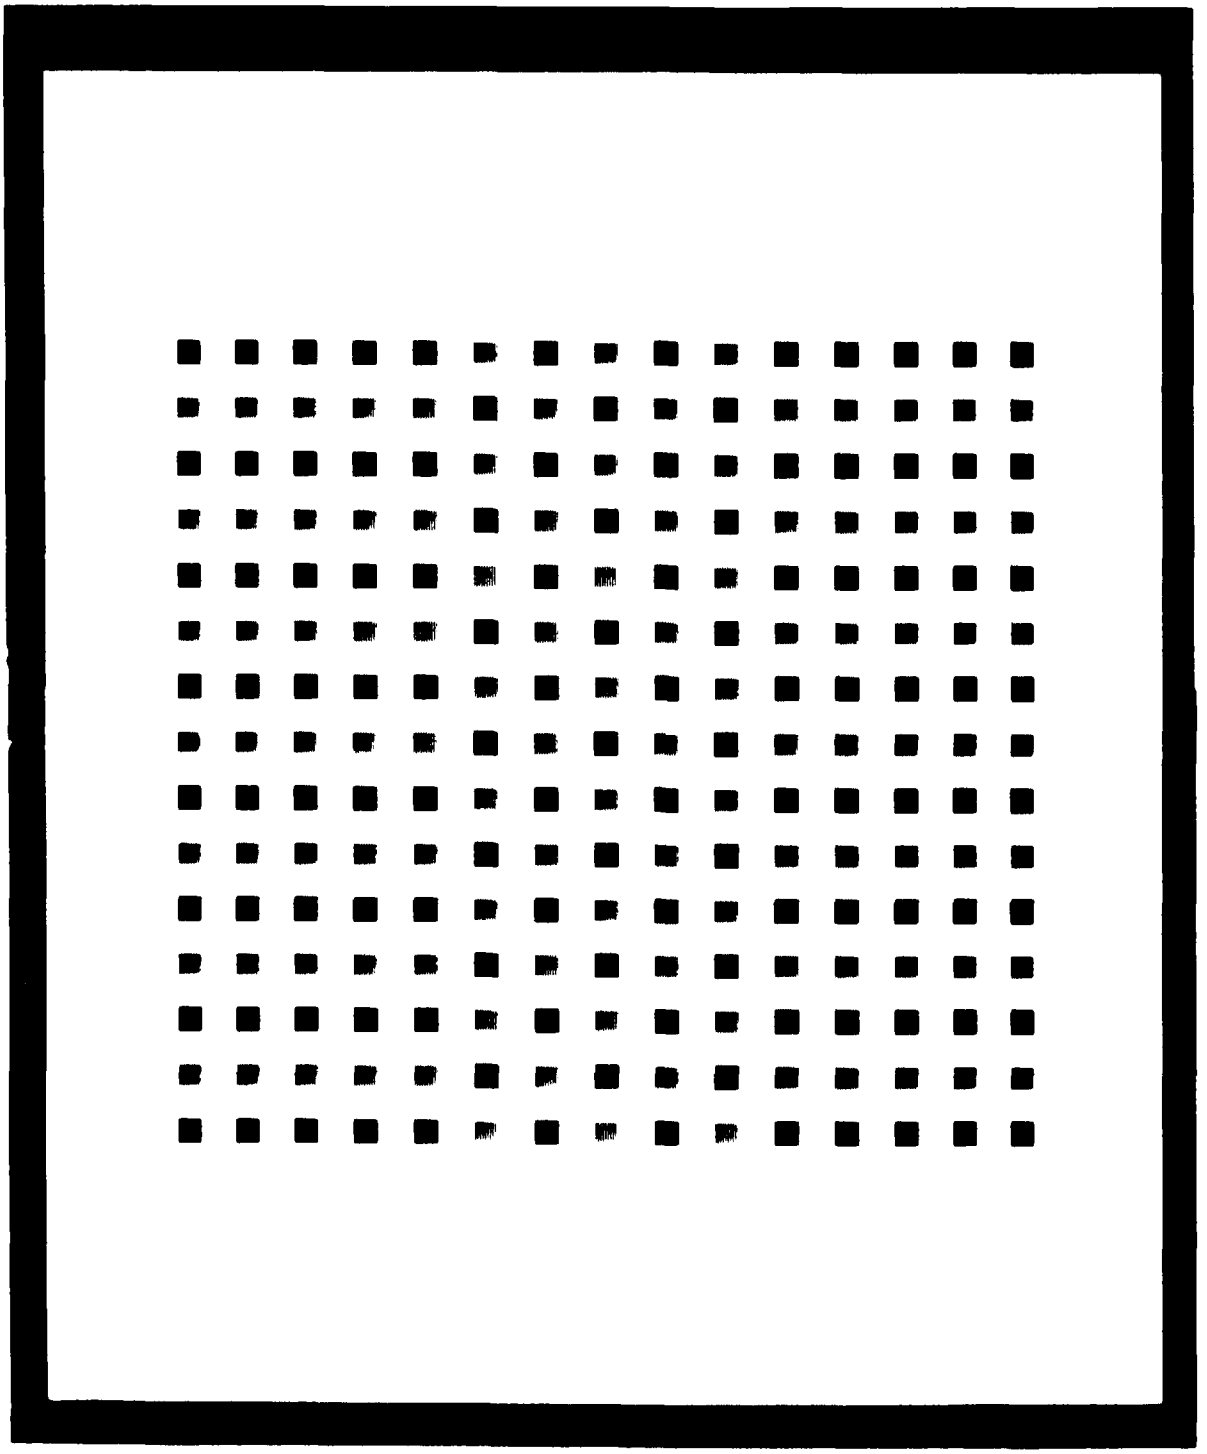

FIGURE 2

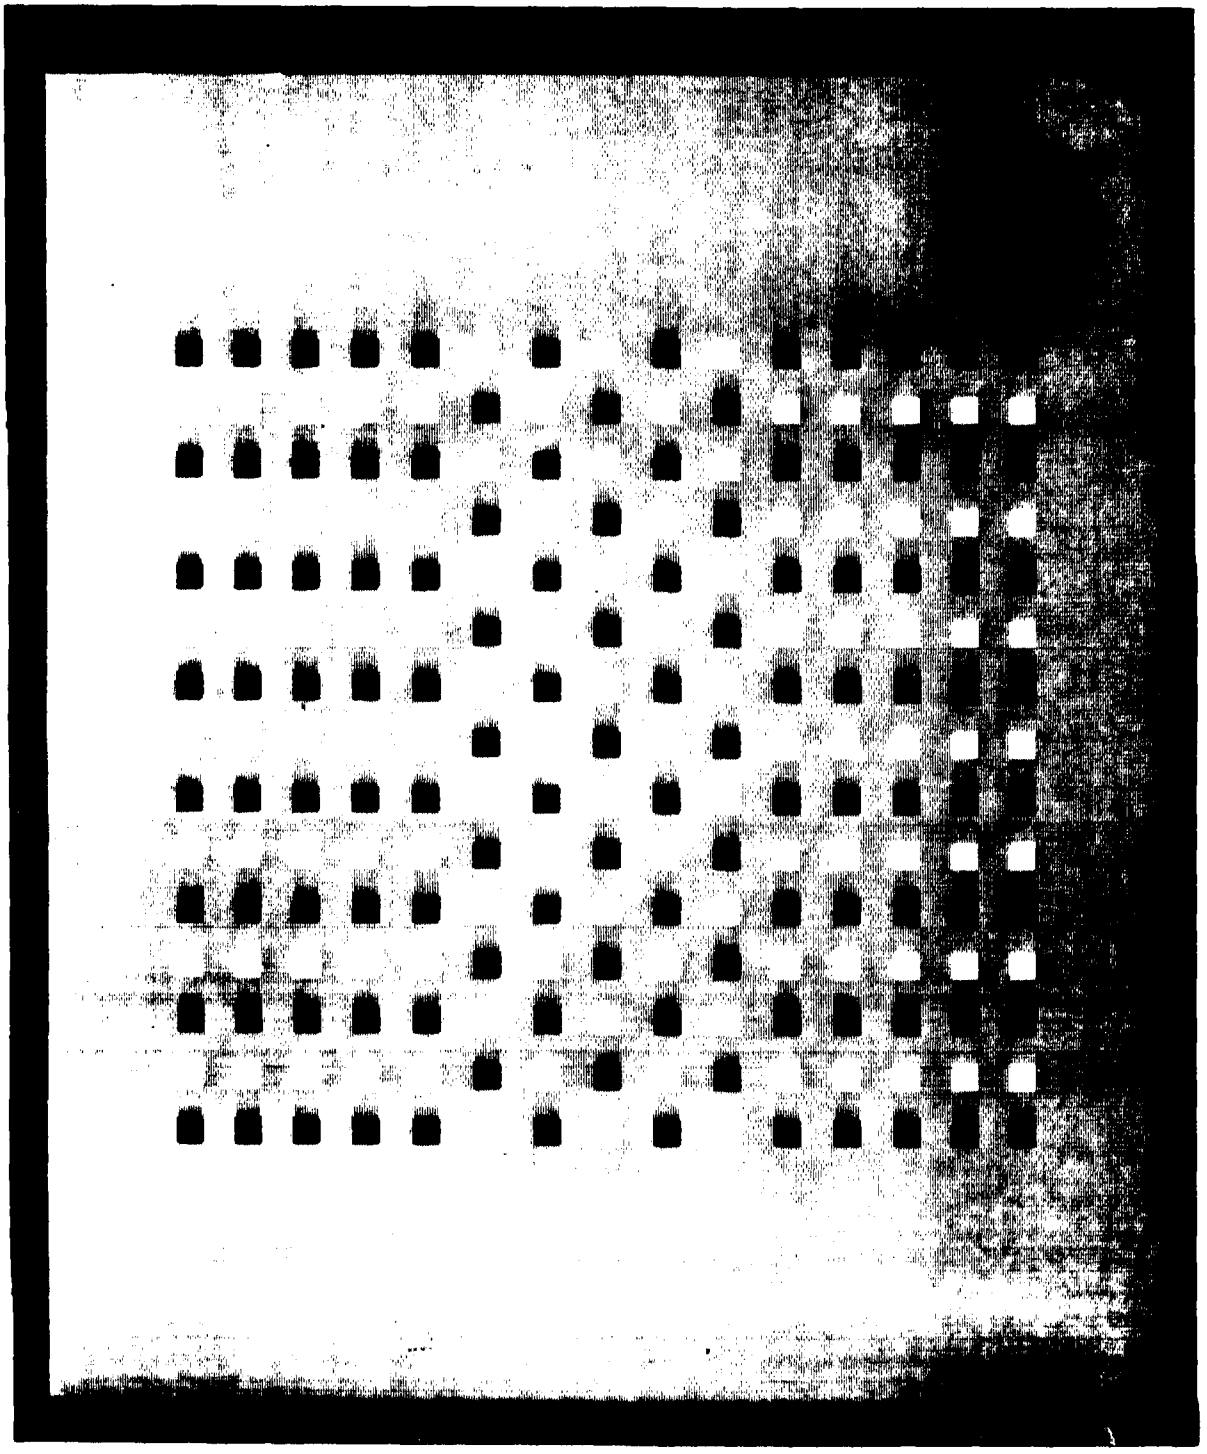

Method

Stimuli. The stimuli were generated by a Symbolics 3600 Lisp machine and displayed on a Tektronix 690 SR color monitor. A stimulus consisted of 15 rows and 15 columns of squares. Each stimulus was divided into three regions, a central section of five rows flanked by a section of five rows above and below. Columns of squares differing in lightness alternated in the top and bottom sections while squares of differing lightness alternated within a column in the middle section. The number of dark and light squares were approximately equal in the three sections. In the top and bottom sections there were 40 dark and 35 light squares. In the middle section there were 37 dark and 38 light squares. Except for Experiment 3, the squares were 10 pixels on a side and the separations between the contours of neighboring squares were 14 pixels. The displays were viewed from a distance of 6 ft. with a pixel subtending 1.08 min of arc.

General Procedure. In Experiments 1 and 4, tripartite segregation was evaluated using a combination of the method of adjustment and a category rating scale. Five subjects rated the degree of segregation of a display when the background intensity was between the intensities of the darker and lighter squares and when the background intensity was above and below the intensities of the squares. The above and below intensities were those in which subjects judged the segregation of a display to be minimal.

A subject was first shown a display with the background intensity between the intensities of the darker (9.2 ft.-L.) and the lighter (17.9 ft.-L.) squares. The subject was told that he would be asked to rate the degree to which the display segmented into three distinct regions on a 5 point scale from 0 to 4. A rating of 0 indicated that the center region did not stand out from the top and bottom regions, i.e. the top, center, and bottom regions appeared to constitute a single pattern. A rating of 4 indicated that the center region stood out strongly from the top and bottom regions. Numbers between 0 and 4 represented intermediate degrees of segregation. The subject was told that he should look at the display normally and not search for the center region.

For 3 subjects, the experimenter then increased the background intensity from its initial value to the maximum value pointing out to the subject that segregation of the display into distinct regions becomes less clear as the background intensity is raised above the intensity of the lighter squares. The subject was told that he would be asked to judge the background intensity at which the segregation of a display into three distinct regions first disappeared or became minimal. He was to do this by having the experimenter raise the background intensity to a value at which the segregation of a display appeared minimal and then to have the experimenter raise and lower the background intensity about this value until he found the lowest background intensity at which the segregation of a display disappeared or became minimal. The subject was told that he would then be asked to rate the degree of tripartite segregation at this background intensity. A subject was given practice determining the background intensity at which the segregation of a display became minimal.

Each subject made 10 ratings of tripartite segregation at the initial between background intensity, 10 judgements of the background intensity at which tripartite segregation became minimal, and 10 ratings of display segregation at the background intensity for which tripartite segregation was judged to be minimal. The between background intensity was varied from tiral to trial and ranged from 11.2 ft.-L. to 16.0 ft.-L. The experimenter then demonstrated that the segregation of a display becomes worse when the intensity of the background is lowered below the intensity of the dark squares. The same procedure was followed as when the intensity of the background was raised. The background intensity was under computer control and could be vried in 128 steps from .02 ft.-L. to 38.0 ft.-L. For two subjects the order was reversed and they were run first with decreaseing the background intensity and then with increasing the background intensity.

In Experiments 2, 3 and 5, the method of category scaling was used to scale the degree of tripartite segregation. At the beginning of a session, each subject was shown displays that segregated strongly, or segregated weakly or not at all. This served to familiarize a subject with the stimuli to be scaled and to acquaint him with the range of variation. The stimuli were presented in a different irregular order to each subject. Except for Experiment 3, the five point rating scale described above was used. The number of subjects and the number of stimulus presenteation varied with each experiment and will be described below.

Subjects. Subjects in the different experiments were drawn from a pool of twelve people. All were naive as to the purpose of the experiment and were paid for their participation.

Results

Experiment 1. Experiment 1 investigated how the background intensity affected tripartite segregation. When the background intensity was between the intensities of the squares, texture segregation occurred. The mean of subjects' segregation ratings was 2.5 with a standard deviation of .68. As shown in Figure 2, raising the background intensity above the intensity of the squares interferes with tripartite segregation. The mean intensity value at which subjects reported tripartite segregation to disappear was 25.9 ft.-L.. The mean of subjects' ratings of tripartite segregation at the disappearance values was .06 with a standard deviation of .09. It is important to note that the lightness difference between the lighter and darker squares is readily discriminable. Texture segregation and the discrimination of lightness depend upon different mechanisms.

Figure 4 shows that lowering the background intensity below the intensity of the darker square also decreases tripartite segregation. (Tripartite segregation may not disappear as completely in Figure 3 as in the laboratory because light adaptation raises the effective intensity of the background.) The mean intensity value at which subjects reported tripartite segregation to disappear was 6.2 ft.-L. The mean of subjects' ratings of tripartite segregation at the disappearance values was .16, with a standard deviation of .30. The ratio of the background intensity to the high-square intensity at which tripartite segregation disappeared when the background

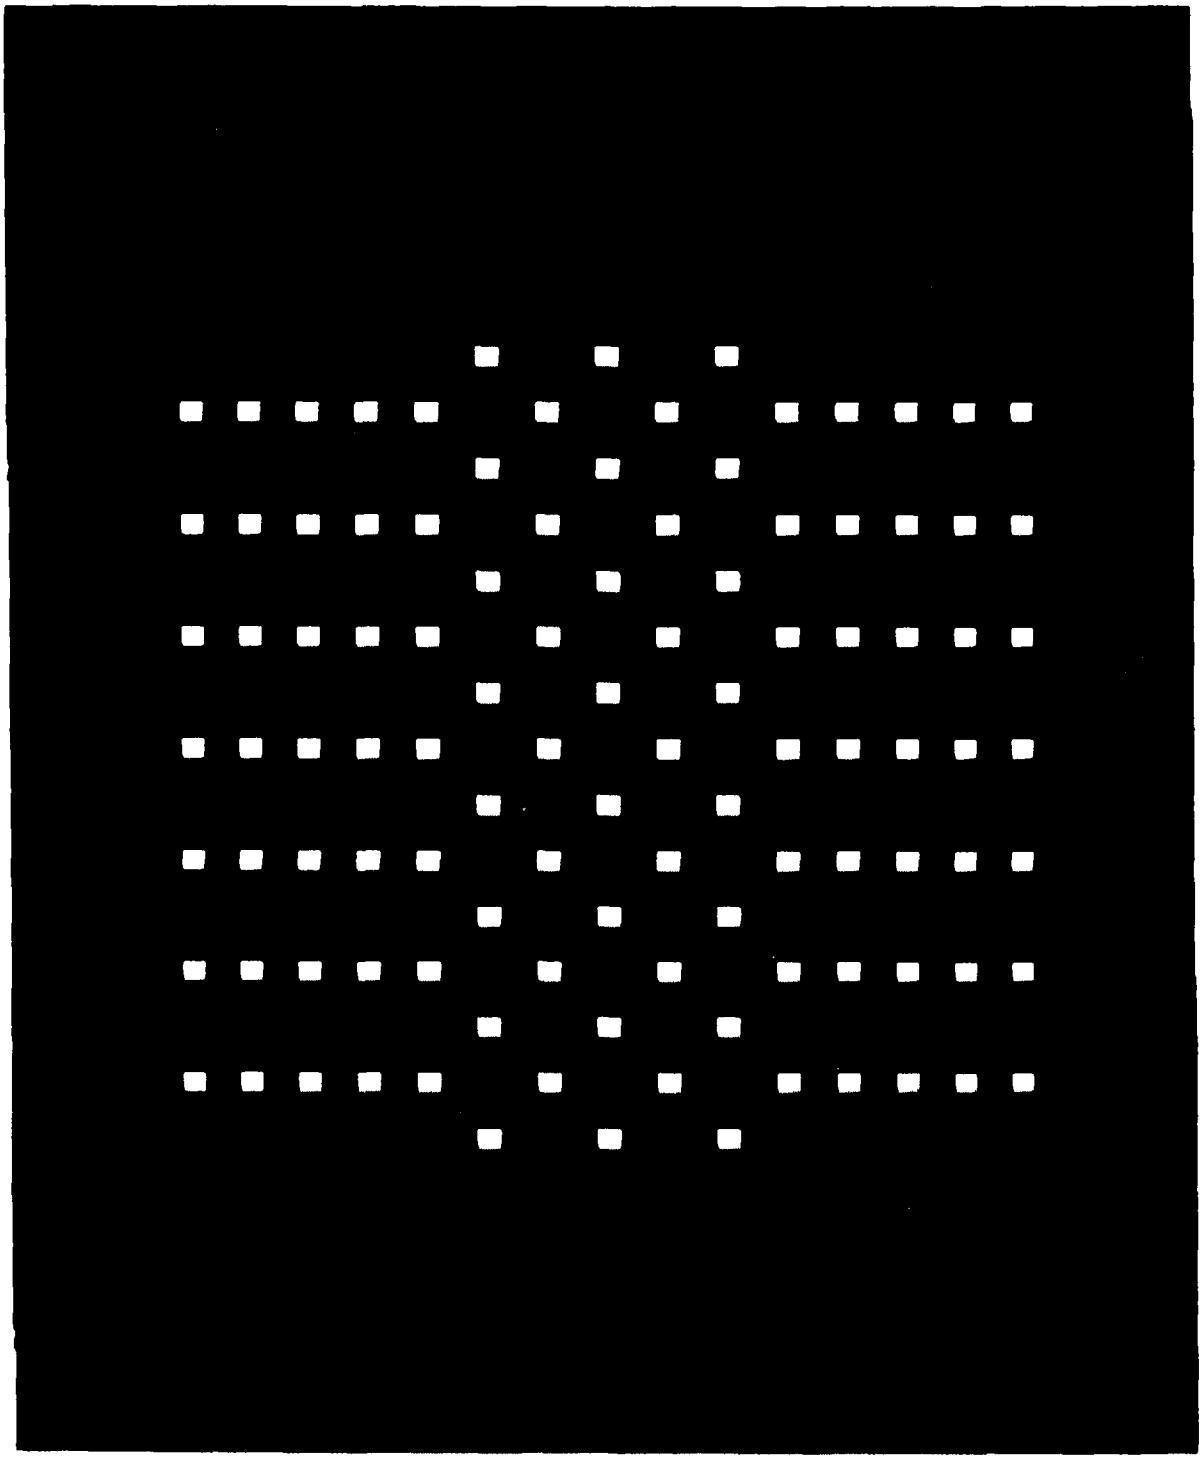

Figure 4

intensity was above was similar to the ratio of the low-square intensity to the background intensity when the background intensity was below, a ratio of approximately 1.5.

Experiment 2. Figure 5 shows that a display consisting of lighter and darker squares on a white background can give a stable tripartite segregation. The lightness difference between the lighter and darker squares, however, has to be large. Our research indicates that the lightness difference has to be between three and five Munsell steps depending on the intensity of the background and the intensities of the lighter and darker squares. Experiment 2 was undertaken to investigate how segregation depends on the background intensity and the intensities of the two squares.

The high-square intensities were set at 12.5, 14.7, 16.6 and 20.8 ft.-L. For each intensity of the high-square, the low-square intensities were varied to give 6 high- to low-square intensity ratios of 1.3, 1.8, 2.6, 5.2, 33 and above 700. At the highest ratio, the low-square intensity was .017 ft.-L., and its ratio to the high-square ranged from 734 to 1224. The background intensity was set in between the high- and low-square intensities and above the high-square intensity. There were four intensity ratios of the background to the high-square when the background intensity was above: 1.2, 1.5, 1.7, and 2.0. An additional condition was run when the background to high-square intensity ratio was 1.2. The high- to low-square intensity ratio was 1.16. There were, thus, 25 experimental conditions with the background intensity between the high- and low-square intensities and 25 experimental conditions with the background intensity above the high- and low-square intensities. A trial consisted of the following sequence: a blank field, a stimulus presented with the background intensity betwen that of the low-square and high-square intensities, a blank field, a stimulus presented with the background intensity above that of the high- and low-square intensities. The stimuli and blank fields were presented for a duration of 1,000 msec. The 25 conditions with the background above were randomly intermixed with the 25 conditions with the background in between. Six subjects served in the experiment, each making 5 ratings of each stimulus.

Figure 6 shows the mean ratings when the background was between the high- and low-square intensities. Rated tripartite segregation was constant and decreased only when the high- to low-square intensity ratios were 1.2 and 1.3. At these intensity ratios it is difficult to see the individual squares, but segregation of the region is still seen. Figure 7 shows the mean ratings when the background was above the high and low square intensities. Tripartite segregation is a function of the ratio of the background intensity to the high-square intensity and of the high-square intensity to the low-square intensity. An analysis of variance revealed that the background to high-square intensity ratio [$F\ (3, 15) = 58.4\ p < .01$], the high- to low-square intensity ratio [$F\ (5, 25) = 40.3\ p < .01$] as well as the interaction between the background to high-square intensity ratio and the high to low-square intensity ratio [$F\ (15, 75) = 5.80\ p < .01$] were significant. When the background to the high-square intensity ratio was 1.2, tripartite segregation greatly improved with increases in the ratio of the high to low-square intensities. As the ratio between the background and high-square intensity is increased, the ratio between the

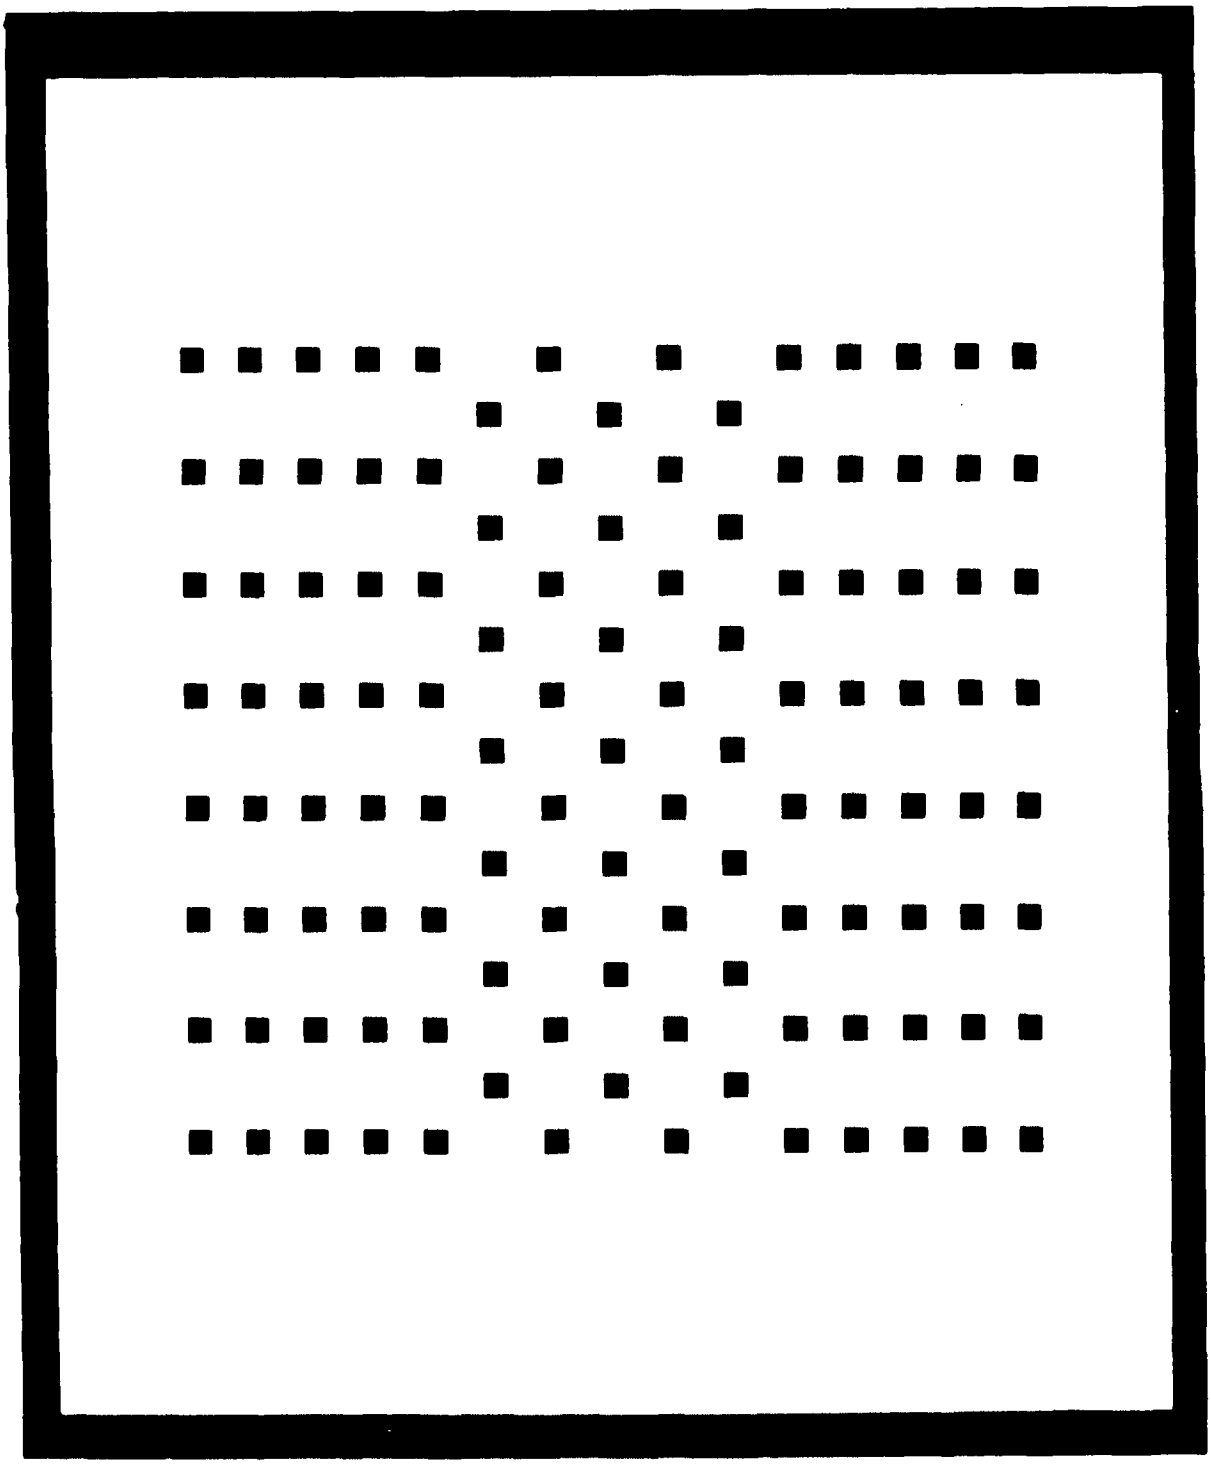

Figure 5

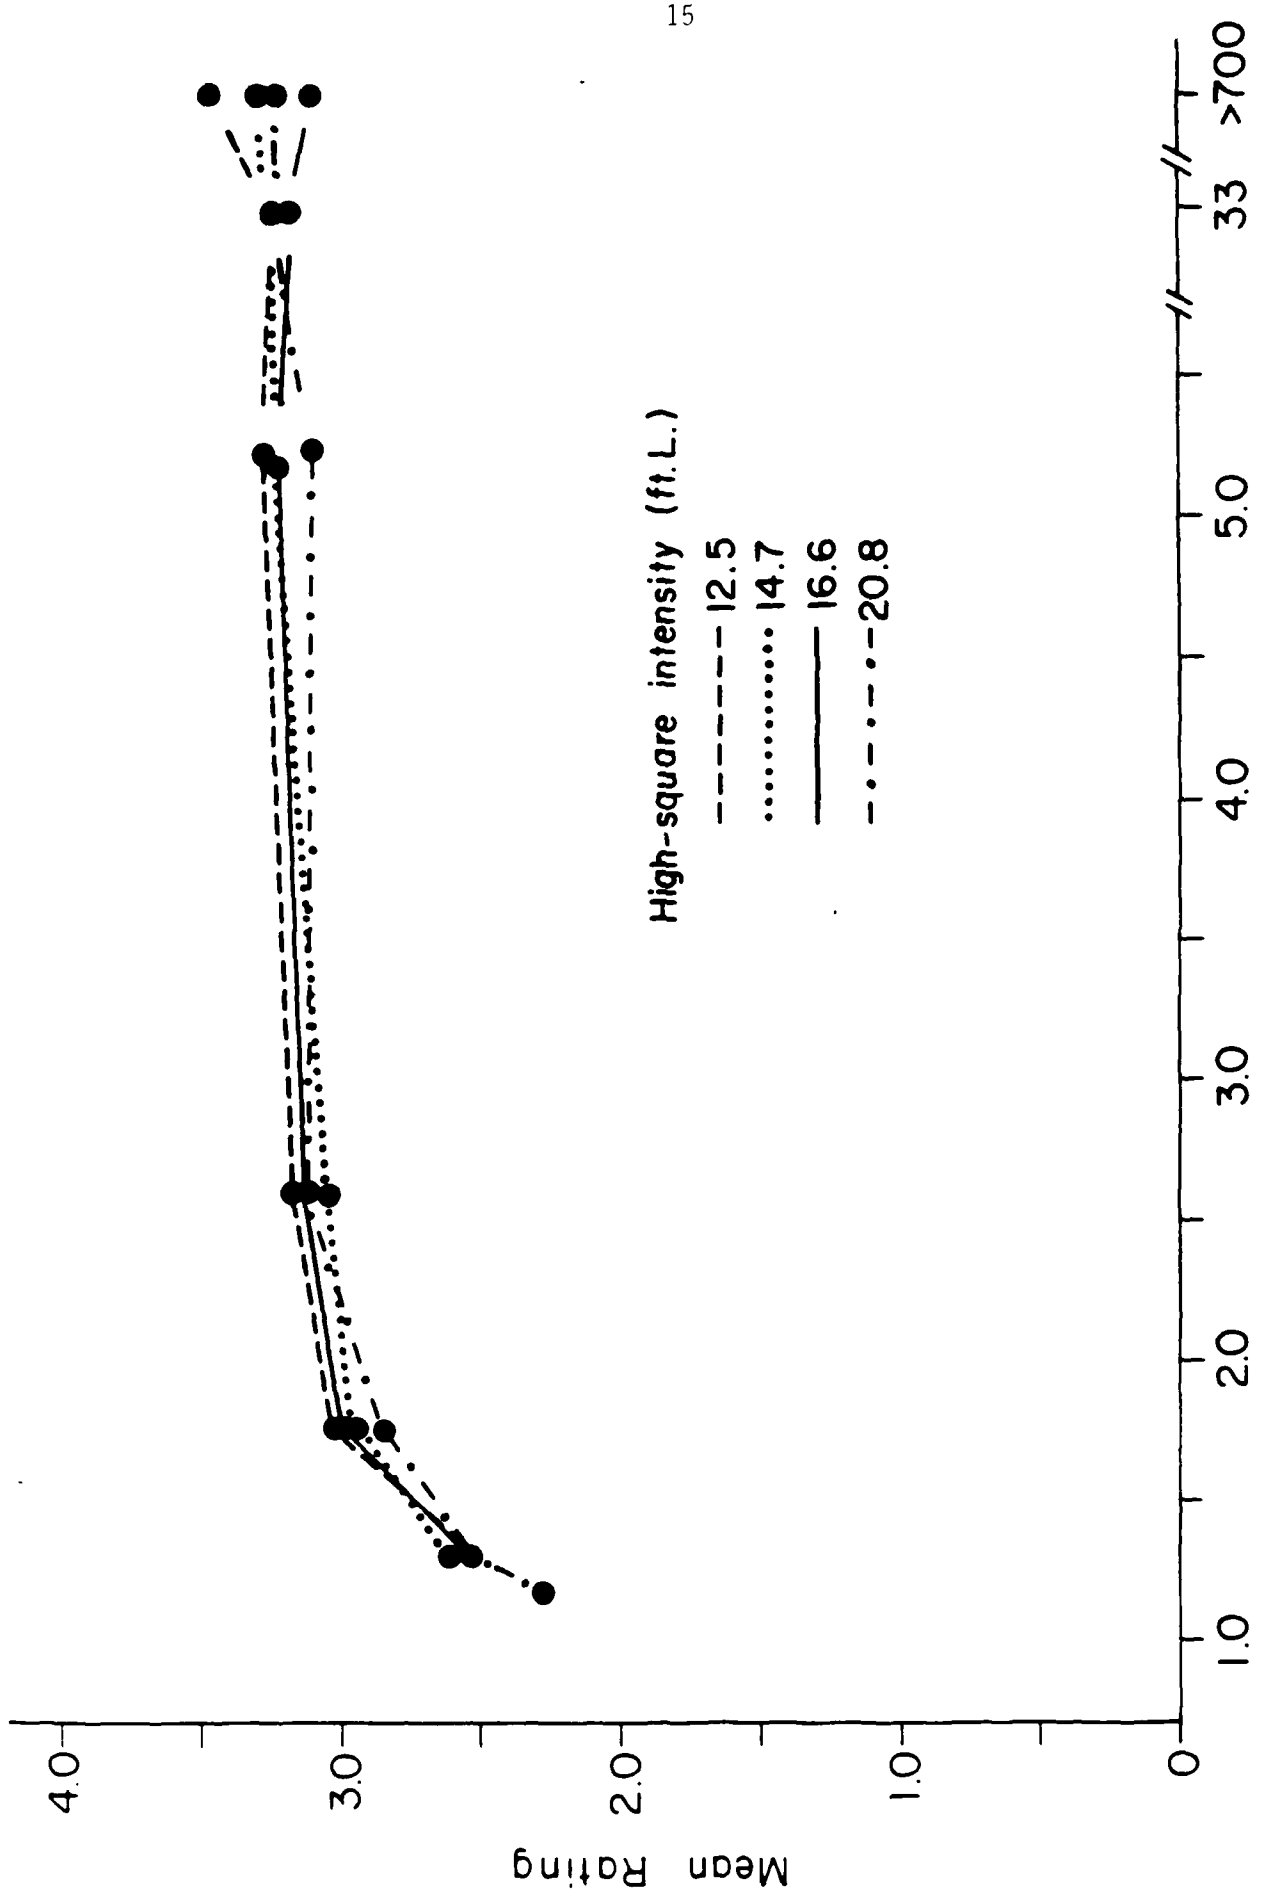
Mean Rating
4.0
3.0
2.0
1.0
0
1.0
2.0
3.0
4.0
5.0
33
>700
Ratio High-/Low-Square Intensity
High-square intensity (ft.L.)
12.5
14.7
16.6
20.8

Figure 6

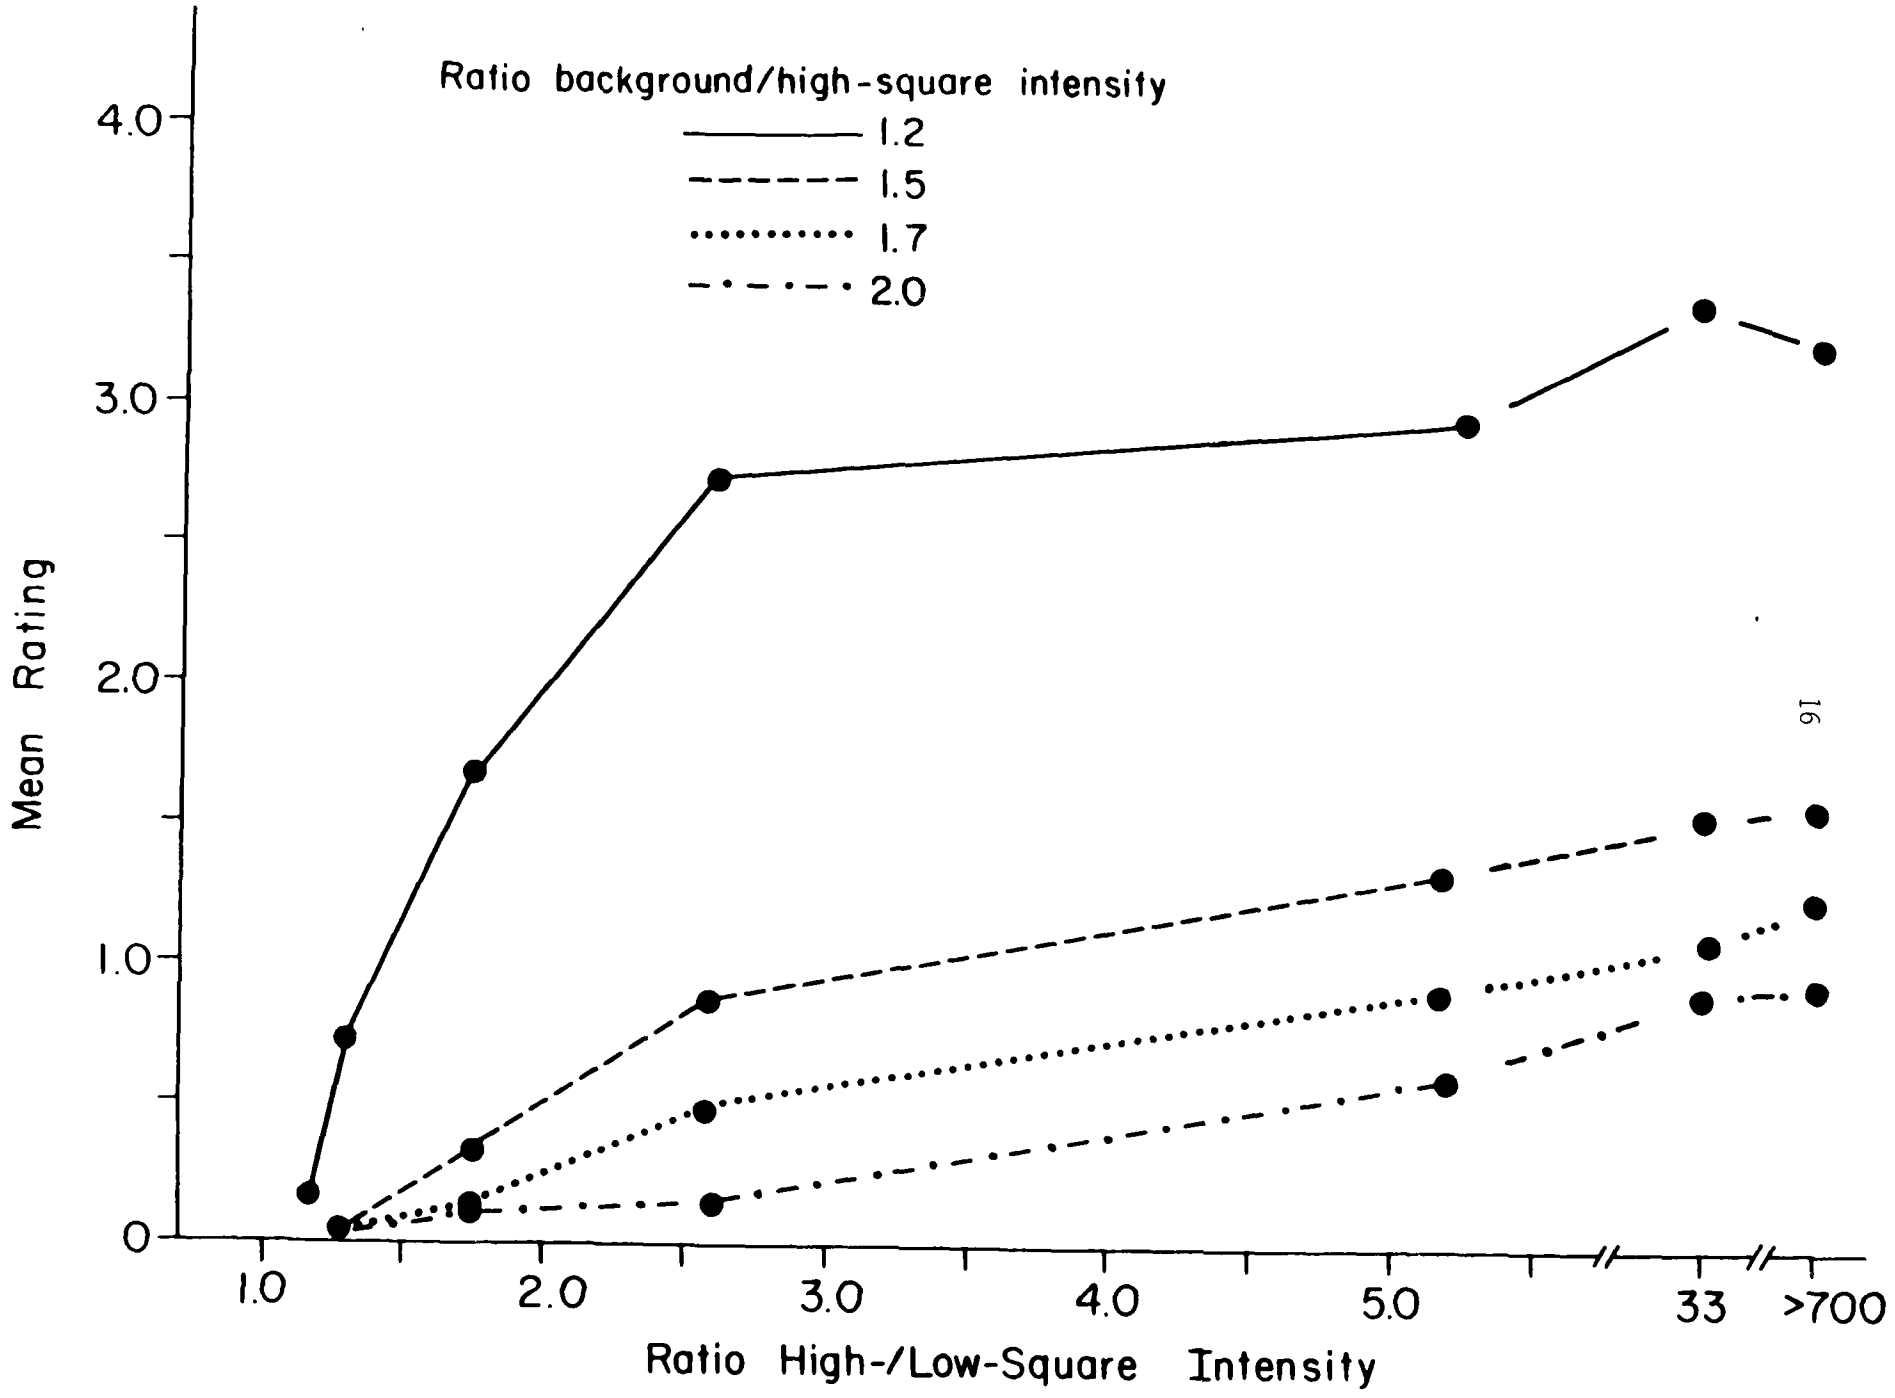
Ratio background/high-square intensity
1.2
1.5
1.7
2.0
Mean Rating
4.0
3.0
2.0
1.0
0
1.0
2.0
3.0
4.0
5.0
33
>700
Ratio High-/Low-Square Intensity

Figure 7

high-to low-square intensity needs to be increased for tripartite segregation to occur. When the ratio of the high- to low-square intensity is above 1.7 or 2.0, tripartite segregation is weak even when the intensity ratio between the high- to low-square intensities is very large. Differences in lightness which are easily seen when individual lightnesses of the squares are compared fail to give tripartite segregation.

Experiment 9. Figure 8 shows that tripartite segregation is strong when the separation betwen the contours of squares is 2 pixels rather than 14 pixels. This is to be expected. What is important is that tripartite segregation now no longer depends on the intensity of the background. Figure 9 shows that tripartite segregation occurs strongly when the background intensity is above that of the squares, and Figure 10 shows that it occurs strongly when the background intensity is below that of the squares. The basis for tripartite segregation is different when the squares are very close. Experiment 3 investigated how the spatial separation of the squares affects tripartite segregation when the background intensity was between and when the background intensity was above that of the squares.

Two conditions in which the background intensity was between the intensities of the squares and three conditions in which the background intensity was above the intensities of the squares were combined with five spatial separations. The between background intensities were set half-way between the intensities of the high-and low-squares. The high-square intensity was set at 26.0 ft.-L. and the low-square intensities were varied to give high- to low-square ratios of 1.4 and 1.8. In the above background conditions, the intensities of the background and high-square were set at 38.0 ft.-L. and 26.0 ft.-L. The low-square intensities were varied to give high- to low-square intensities of 1.4, 1.8, and 1040. The three ratios with the background itensity above were judged to produce threshold, weak, and strong segregation with a contour to contour separation of 14 pixels. The contour to contour spatial separations of the squares were 2, 6, 10, 14, and 18 pixels. The stimulus with a high to low-square intensity ratio of 1.4, the background intensity between, and with a 14 pixel contour to contour spatial separation served as a standard and was assigned a value of 5. It is shown as a triangle in Figure 11. Subjects were instructed to rate the degree of segregation on an 11 point scale from 0 to 10. The standard was presented on each trial. A trial consisted of the followin sequence: a blank field, standard stimulus, a blank field, and a comparison stimulus. The durations of the blank fields and stimuli were 1000 msec. Ten subjects served in the experiment making five ratings of each stimulus.

Figure 11 shows the means of subjects' tripartite segregation ratings. Tripartite segregation was consistently rated to be better when the background intensity was between than when the background intensity was above. Separate analyses of variance were conducted with the background intensity above and with the background intensity between. When the background intensity was between, spatial separation [$F(4, 36) = 136.3$ $p < .01$], and the ratio of the high- to low-square intensity [$F(1, 9) = 17.1$ $p < .01$] were significant. The spatial separation by the ratio of high- to low-square intensity was not significant [$F(4, 36) = 2.4$ $p < .05$].

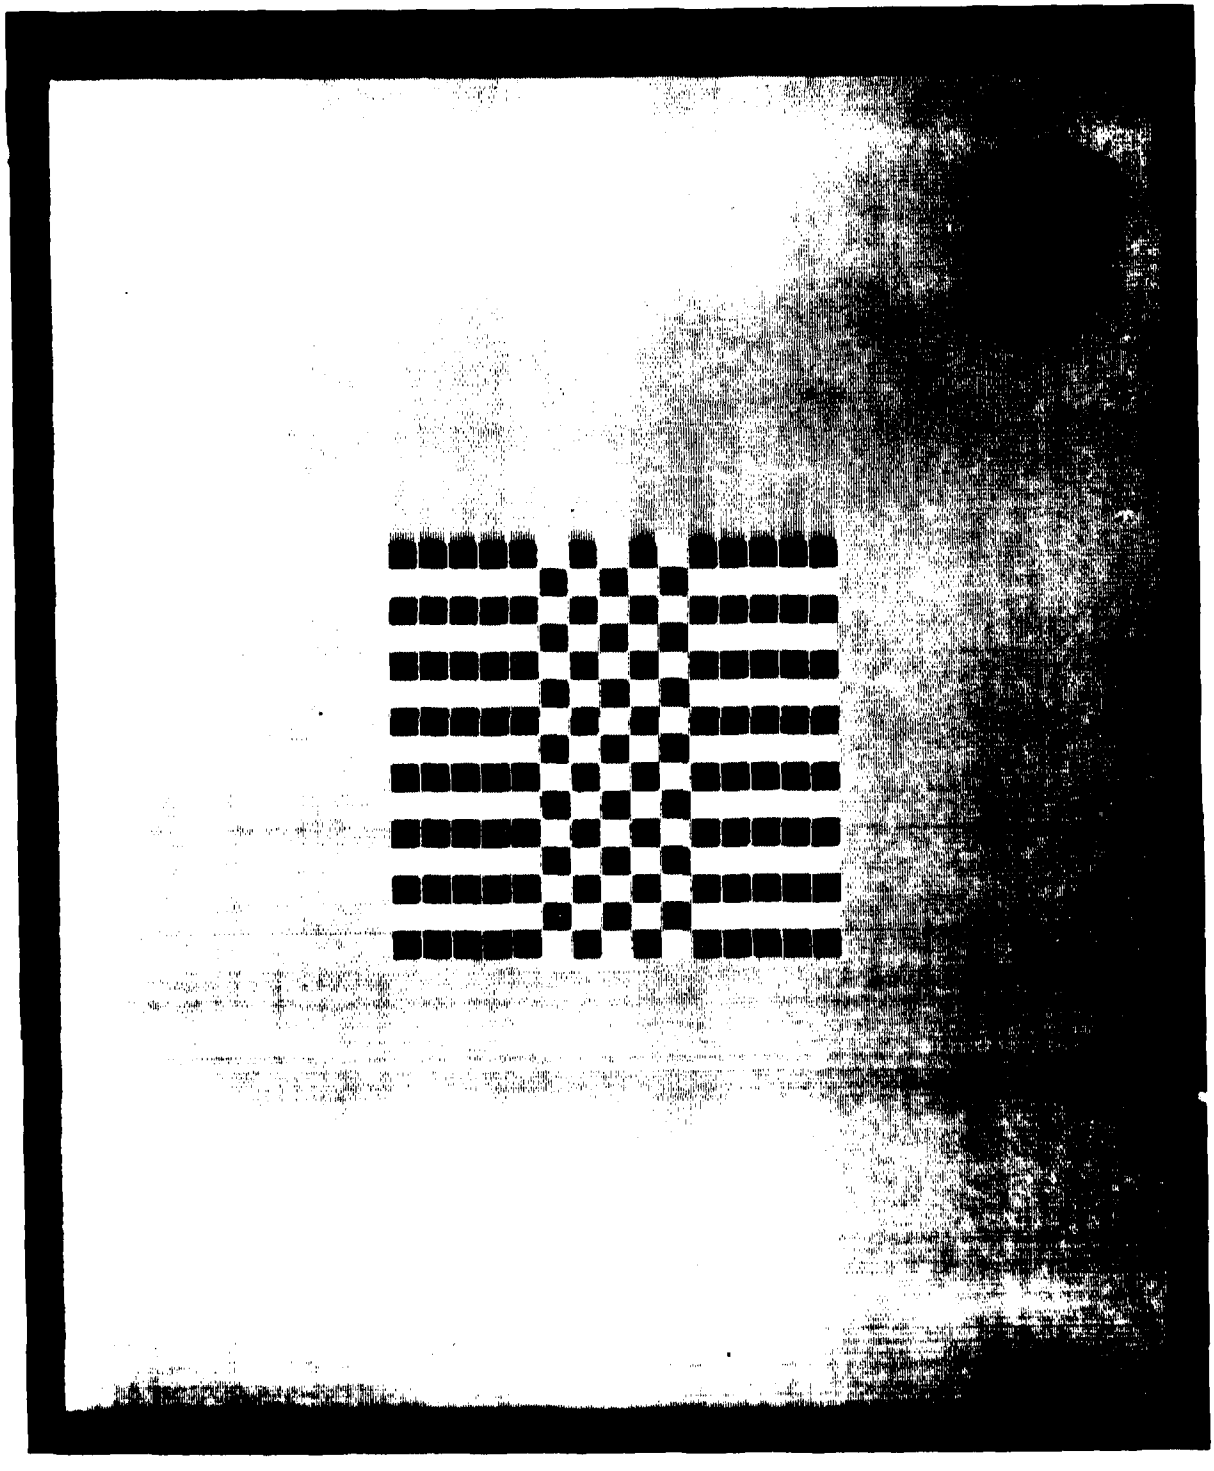

Figure 8

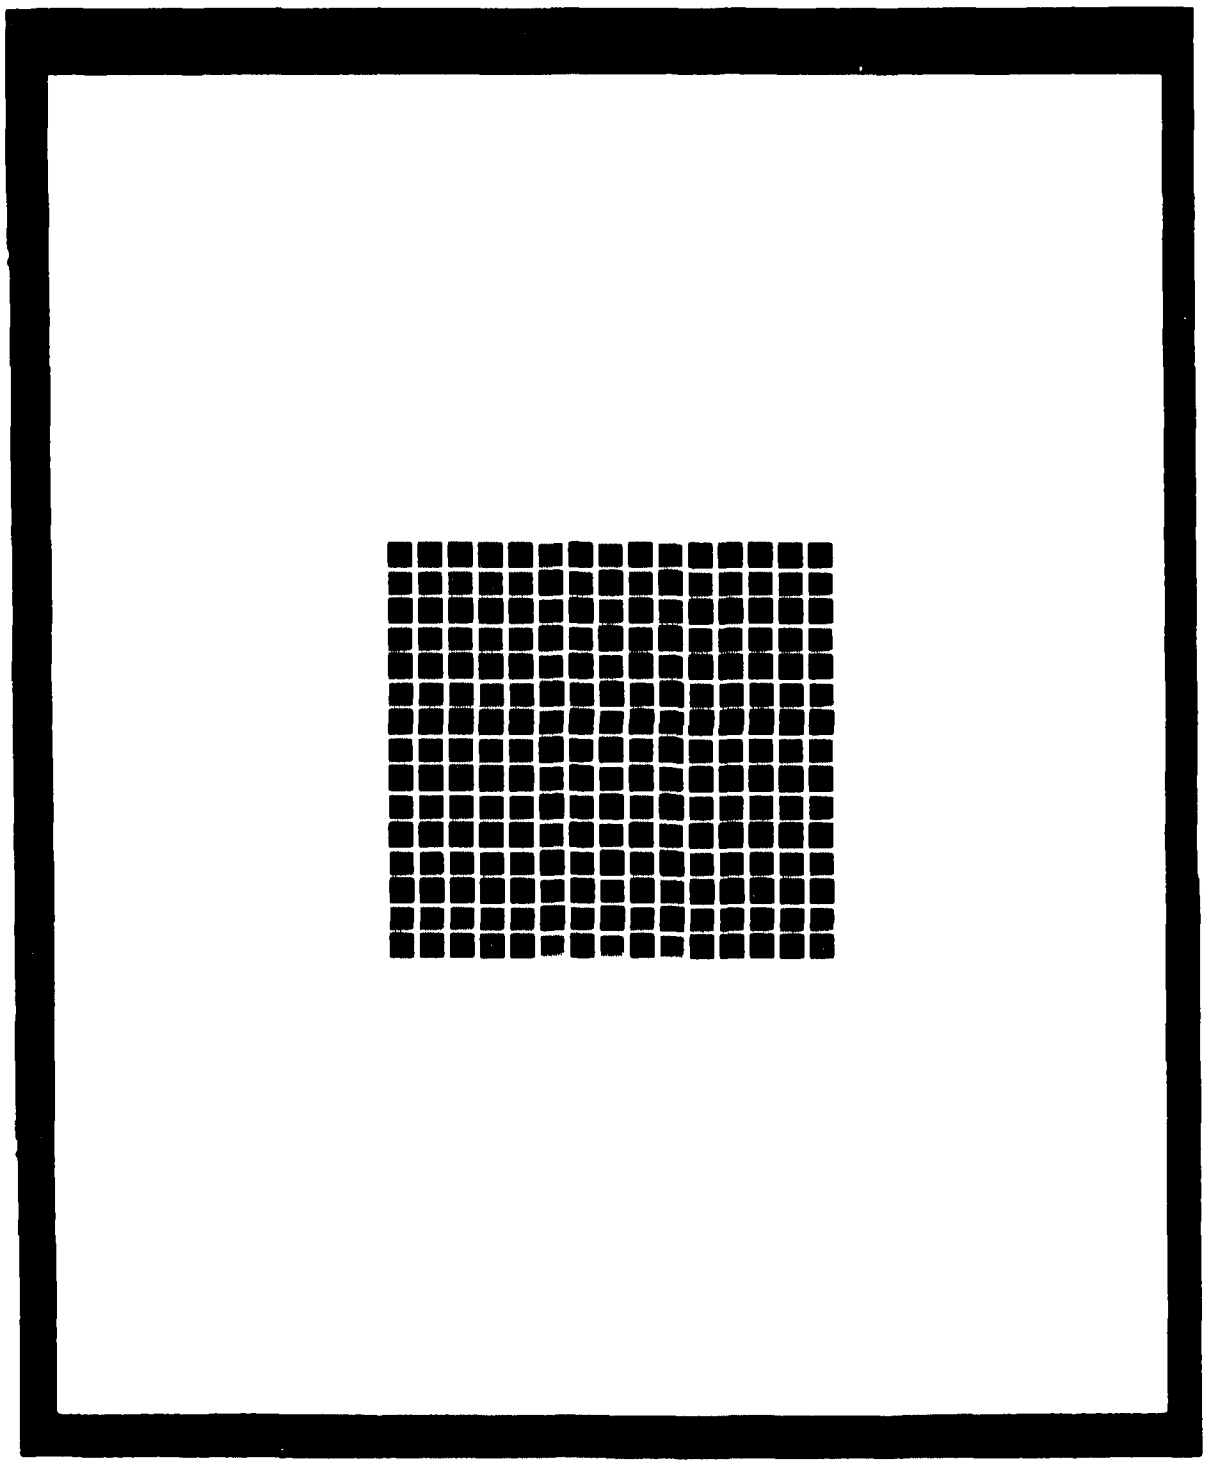

Figure 9

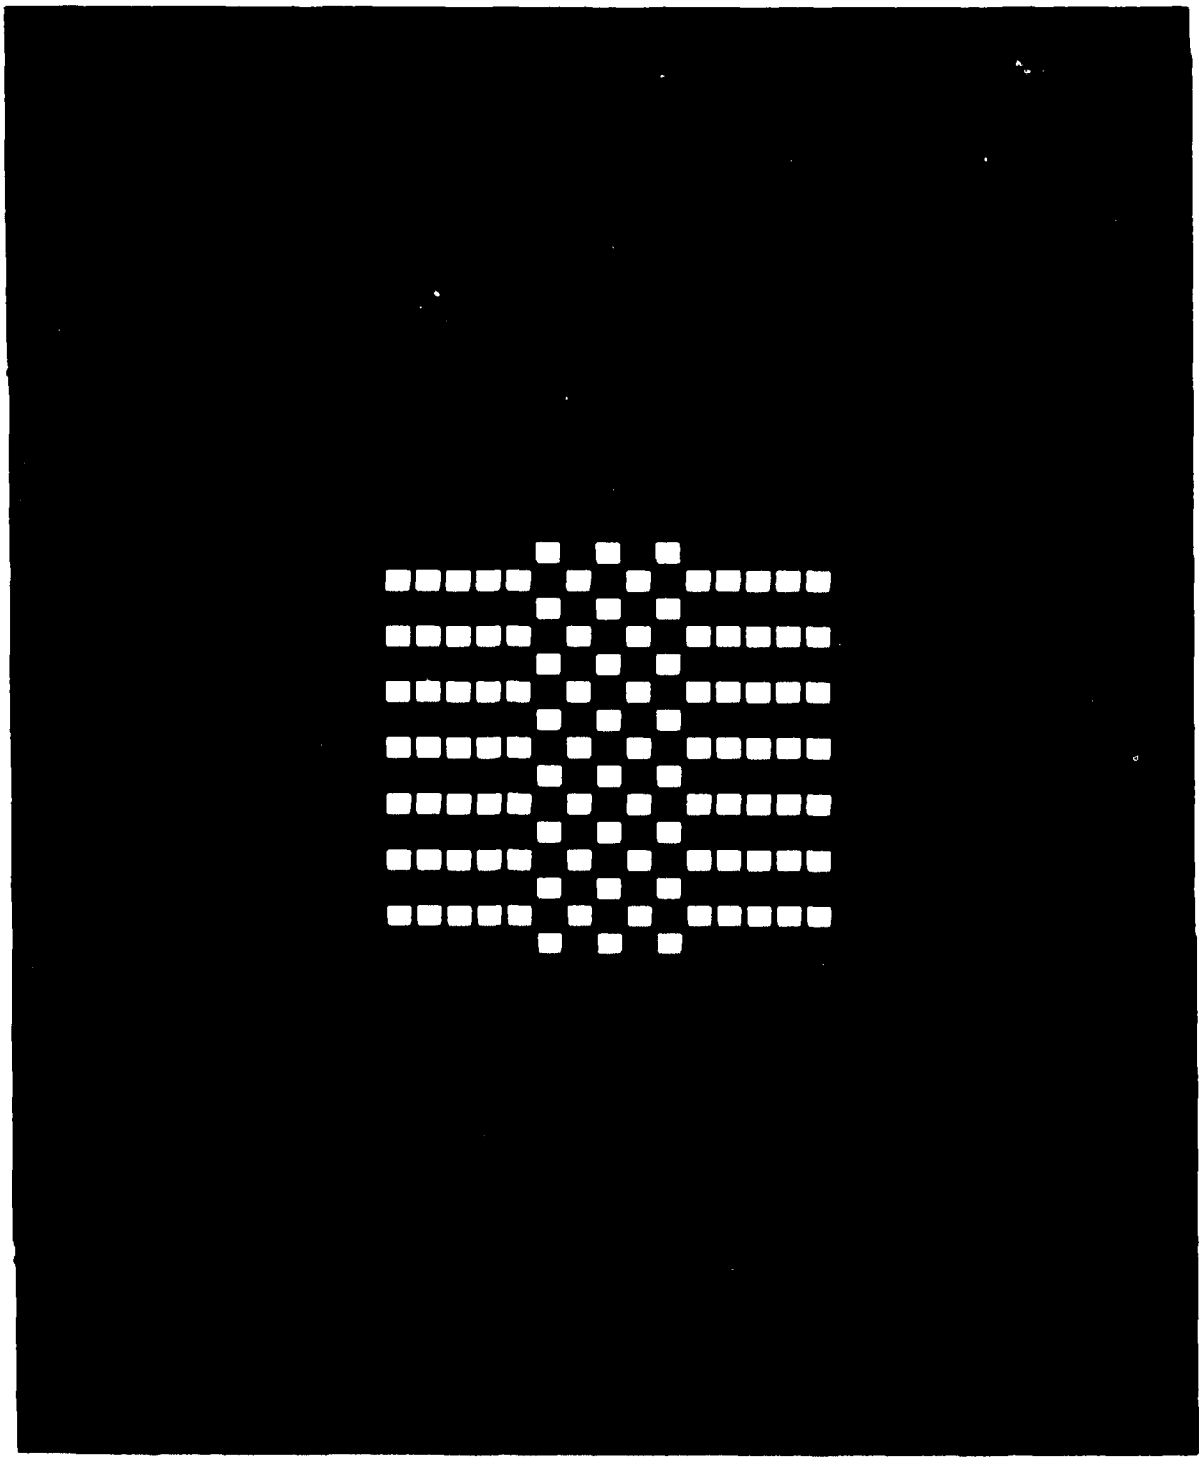

Figure 10

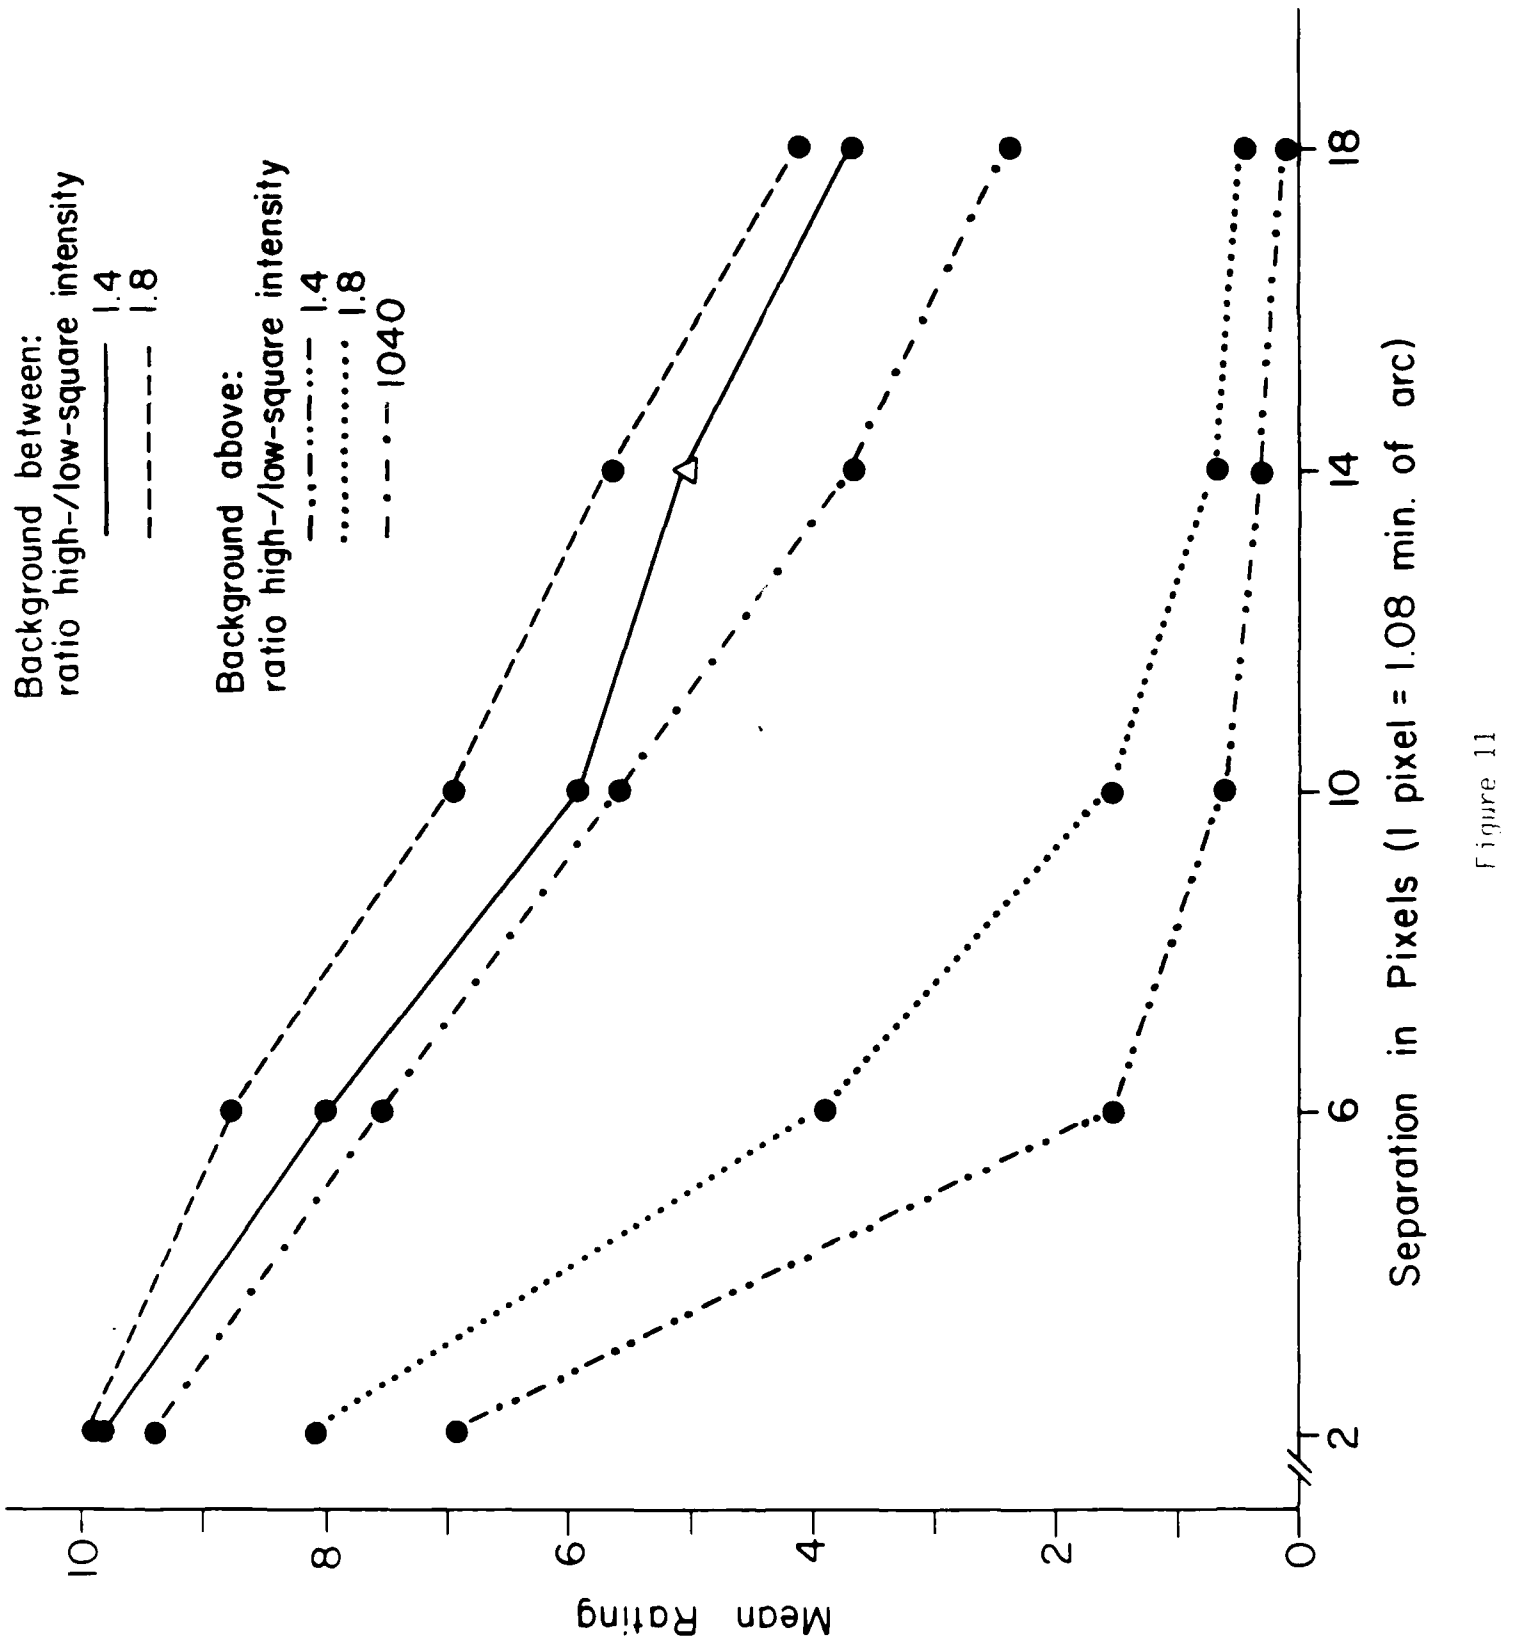
Background between:
ratio high-/low-square intensity
1.4
1.8
Background above:
ratio high-/low-square intensity
1.4
1.8
1040
Mean Rating
0
2
4
6
8
10
2
6
10
14
18
Separation in Pixels (1 pixel = 1.08 min. of arc)

Figure 11

When the background intenity was above, spatial separation [F (4, 36) = 187.9 p < .01]. the ratio of the high- to low-square int ·s y [F (2,18) = 169.4 p < .01], and the separation by high- to low-square intensity were significant [F (8, 72) = 13.8 p < .01]. The interaction reflects the fact that the rated segregation when the ratio of the high- to low-square intensity was 1040 decreased gradually with spatial separation similar to that which occurs when the background intensity was between. When the ratios of the high- to low-square intensities were 1.4 and 1.8, rated tripartite segregation was strong only with a 2 pixel separation and decreased strongly with increasing spatial separation.

Experiment 4. The procedure in Experiment 4 was the same as in Experiment 1. Subjects were instructed to report when the tripartite segregation became minimal or disappeared when the background intensity was raised from between to above and below that of red and green squares. The intensity of the red squares was 6.4 ft.-L., and of the green squares was 18.0 ft.-L. The background intensity ranged from 31.5 to 16.0 ft.-L. The mean of subjects' tripartite segregation ratings was 2.4 with a standard deviation of .60. The mean intensity of the background for which tripartite segregation was reported to disappear when it was raised above the intensities of the squares was 25.8 ft.-L. Subjects' mean ratings of tripartite segregation was .12 with a standard deviation of .27. The mean intensity of the background at which tripartite segregation was reported to disappear when the background intensity was decreased was 3.7 ft.-L. The mean of subjects' tripartite segregation ratings was .2. One can see the difference betwen the red and green squares, but this is not sufficient to segment the stimulus into distinct regions. One has to search out the vertical columns of red and green squares present in the top and bottom regions and absent in the middle region. Though red and green are opponent colors (that is, coded as opposites by the visual system), a white background fails to segregate red and green squares in the same way as a gray background segregates squares whose intensity is greater than the background from the squares whose intensity is less than the background.

Following the judgments with the red and green displays, the subject scaled the degree of segregation of red and blue squares on a purple background and of yellow and blue squares on an achromatic background. Figure 12 shows red and blue squares with the background intensity between. Tripartite segregation is strong. Figure 13 shows the same red and blue squares with the background intensity above. Tripartite segregation has been greatly lessened. The order of presenting the red and blue square displays, and the yellow and blue square displays were alternated. The background intensity was set between and above the intensities of the two squares. The red squares were 2.8 ft.-L. and the blue squares 1.1 ft.-L. When the intensity of the purple background was in betwen 1.3 ft.-L., the mean segregation ratings was 2.6. When the purple background intensity was raised to 5.5 ft.-L., the mean rating for segregation was .6. Red and blue squares on a purple background whose intensity is greater than that of the red and blue squares fails to produce segregaton, though phenomenally red is to one side of purple and blue is to the other side of purple. Similar results were obtained with the yellow and blue squares. The intensity of the blue square was 2.7 ft.-L. and of the yellow squres 16.2 ft.-L.

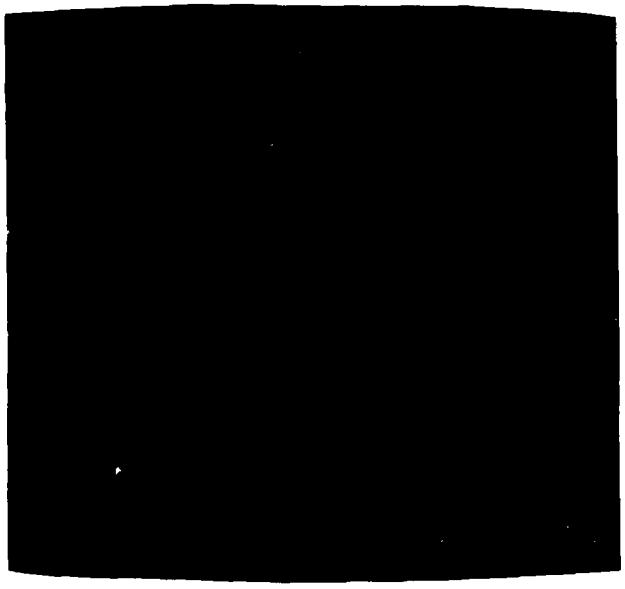

Figure 12

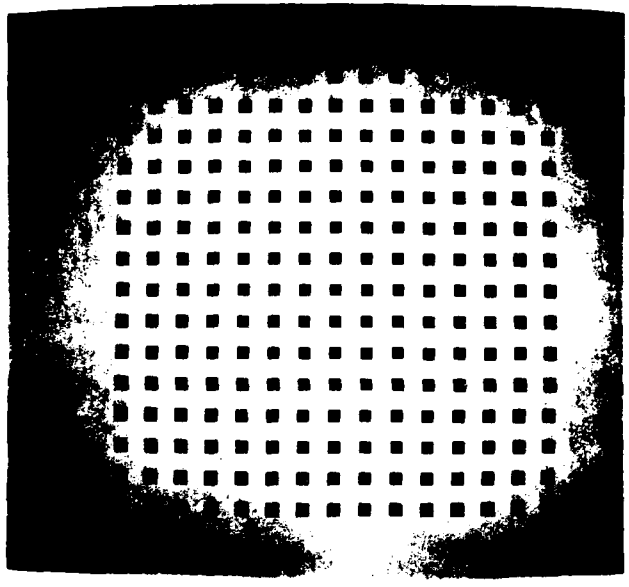

Figure 13

The intensity of the background when between was 11.9 ft.-L. and when above 32.0 ft.-L. The mean segregation rating when the intensity of the background was betwen that of the yellow and blue squares was 3.2 and when the intensity of the background was above that of the yellow and blue squares .80. Thus, tripartite segregation depends on a qualitative difference, sign of the contrast in the luminance system, and not on a qualitative difference in the hue system or in a phenomenological classification. As with achromatic squares, strong tripartite segregation occurs with the background intensity above when the yellow and blue or red and blue squares are separated by 2 rather than by 14 pixels.

Experiment 5. Ordinarily, similarity of hue is effective in producing grouping (Treisman, 1982). Why do not the squares in the top and bottom regions of the tripartite display group into columns on the basis of hue? Experiment 5 was undertaken to determine whether tripartite segregation would occur when the red and blue squares and the background are made isoluminant. The procedure was the same as in Experiment 1. Five subjects were asked to report when segregation disappeared or was minimal when the background was increased in intensity. Each subject made ten judgments. The red and blue squares and background were set at 1.4 ft.-L. and at 2.8 ft.-L. The order in which the two intensities were presented were alternated. The mean intensity at which segregation disappeared when the red and blue squares were 1.4 ft.-L. was 3.4 ft.-L. and when the red and blue squares were 2.8 ft.-L. was 6.3 ft.-L. When the luminances of the achromatic background and the red and blue square were equated, at either 1.4 ft.-L. or 2.8 ft.-L. tripartite segregation occurred. The mean rating when the squares and background were 1.4 ft.-L. was 2.1 with a standard deviation of .73, and when the squares and background were 2.8 ft.-L. 2.3 with a standard deviation of .97. These means do not differ significantly. The mean tripartite segregation ratings when the background intensities were raised above that of the squares was 0. Hue differences can give tripartite segregation but that tripartite segregation based on hue differences are overridden by the similarity in contrast when the background is above or below that of the squares.

2.3 Tripartite Experiments 6-8: Unequal Size Squares

Experiments 6 through 8 studied tripartite segregation resulting from size differences. The squares were the same in lightness and differed in size.

Method

The procedure was the same as in Experiment 1 through 5. Columns of large and small squares alternated in the top and bottom regions. Large and small squares alternated within a column in the middle region. Except for Experiment 8, the large squares were 16 pixels on a side, the small squares 8 pixels on a side, and the separation between the contours of neighboring squares 12 pixels.

Experiment 6. Figure 14 shows that a size difference produces tripartite segregation. The background intensity was 32.0 ft.-L. and the intensities of the large and small squares were 13.8 ft.-

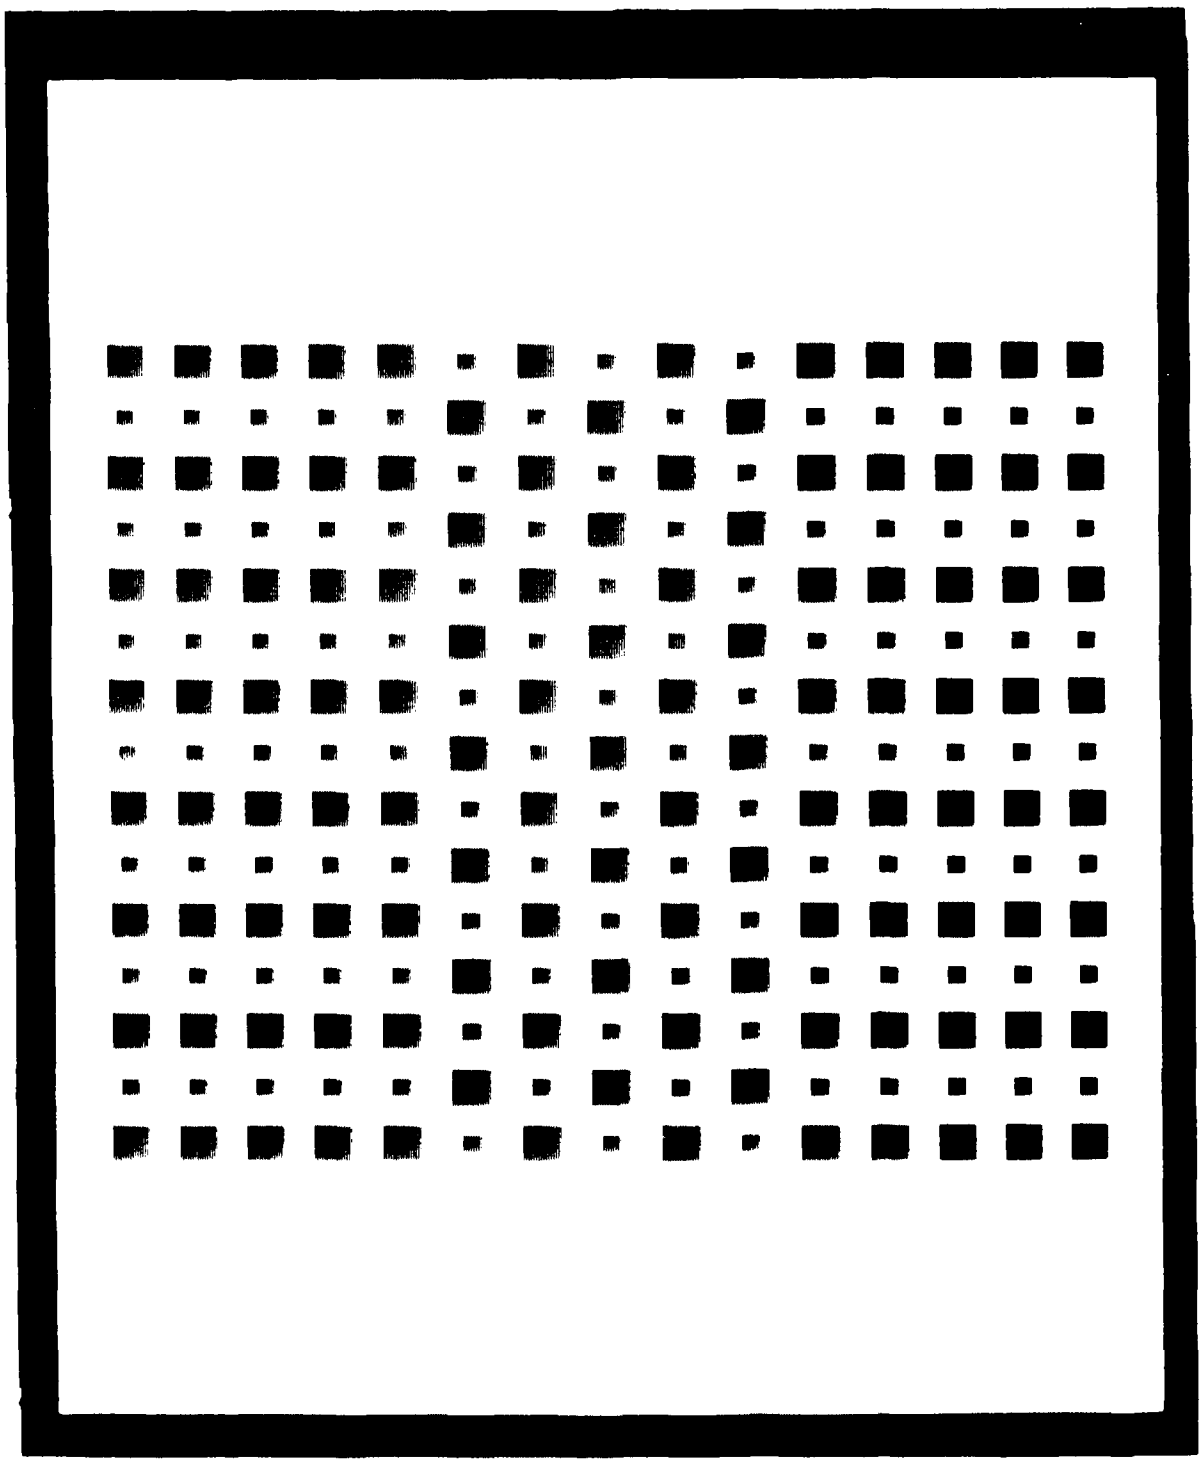

Figure 14

L. The mean segregation rating was 2.6 with a standard deviation of .71. Figure 15 shows the same display when the small squares are made darker. Tripartite segregation is lessened. The mean intensity value at which subjects reported tripartite segregation to disappear was 4.7 ft.-L. The mean of subjects ratings of tripartite segregation was .58 with a standard deviation of .45. Increasing the intensity of the large squares and thereby decreasing their contrast also decreased tripartite segregation. The mean intensity value at which subjects reported tripartite segregation to disappear was 23.0 ft.-L. The mean of subjects rating of tripartite segregation was .24 with a standard deviation of .43. Similar results were obtained with a black background. Figure 16 shows that tripartite segregation occurs when the large and small squares are of equal intensity with a black background. The background intensity was .02 ft.-L. and the intensities of the large and small squares were 5.5 ft.-L. The mean of the five subject segregation rating was 2.5 with a standard deviation of .51. Tripartite segregation is lessened. The mean intensity value of the small squares at which subjects reported tripartite segregation to be minimal was 14.6 ft.-L. The mean of subjects ratings of tripartite segregation at this intensity value was .4 with a standard deviation of .55. Tripartite segregation is also decreased when the large squares are made darker. This is shown in Figure 17. The background was again set at .02 ft.-L. and the small and large squares initially were 5.5 ft.-L. The mean segregation rating at these values was 2.1 with a standard deviation of .48. The mean intensity value at which subjects reported tripartite segregation to be minimal was 1.1 ft.-L. The mean of subjects ratings of tripartite segregation at this intensity value of the large squares was .1 with a standard deviation of .3.

Experiment 7. In Experiment 7 there were two sizes of the small squares. The large square was 16 pixels on a side; the small squares were either 8 pixels on a side or 6 pixels on a side. The contour to contour separation between the squares was either 12 pixels or 2 pixels. The intensity of the small squares was set at 21.0 ft.-L. The intensities of the large squares were varied from .14 to 38.0 ft.-L. For the stimuli with a 12 pixel separation, there were 5 intensities of the large square; .14, .61, 1.6, 4.9, and 38.0 ft.-L. For the stimuli with a 2 pixel separation, there were 3 intensities of the large squares: 1.6, 4.9, and 38.0 ft.-L. The stimuli were paired according to both the size of the small square and the spacing between the squares. In a given trial, for example, a display with an 8 pixel small square and a 12 pixel spacing was presented with a display with a 6 pixel small squares and a 12 pixel spacing. A trial consisted of the following sequence. A blank field, the first stimulus, a blank field, and a second stimulus. The stimuli were presented for 2000 msec and the blank fields for 1000 msec. The order of presenting displays with an 8 pixel square and a 6 pixel small square were counter balanced. Subject made ten scaling judgments. The intensity of the blank field was 38.0 ft.-L. and the background intensity was .02 ft.-L.

Figure 18 presents the results. Tripartite segregation was strong with a 2 pixel separation and varied relatively little as a function of the luminance of the large squares. An analysis of variance revealed that luminance [$F(2, 18) = 17.8$; $p < .01$] and size [$F(10, 4) = 20.0$; $p < .01$] as well as the luminance by size interaction [$F(2, 18) = 6.4$; $p < .01$] were significant.

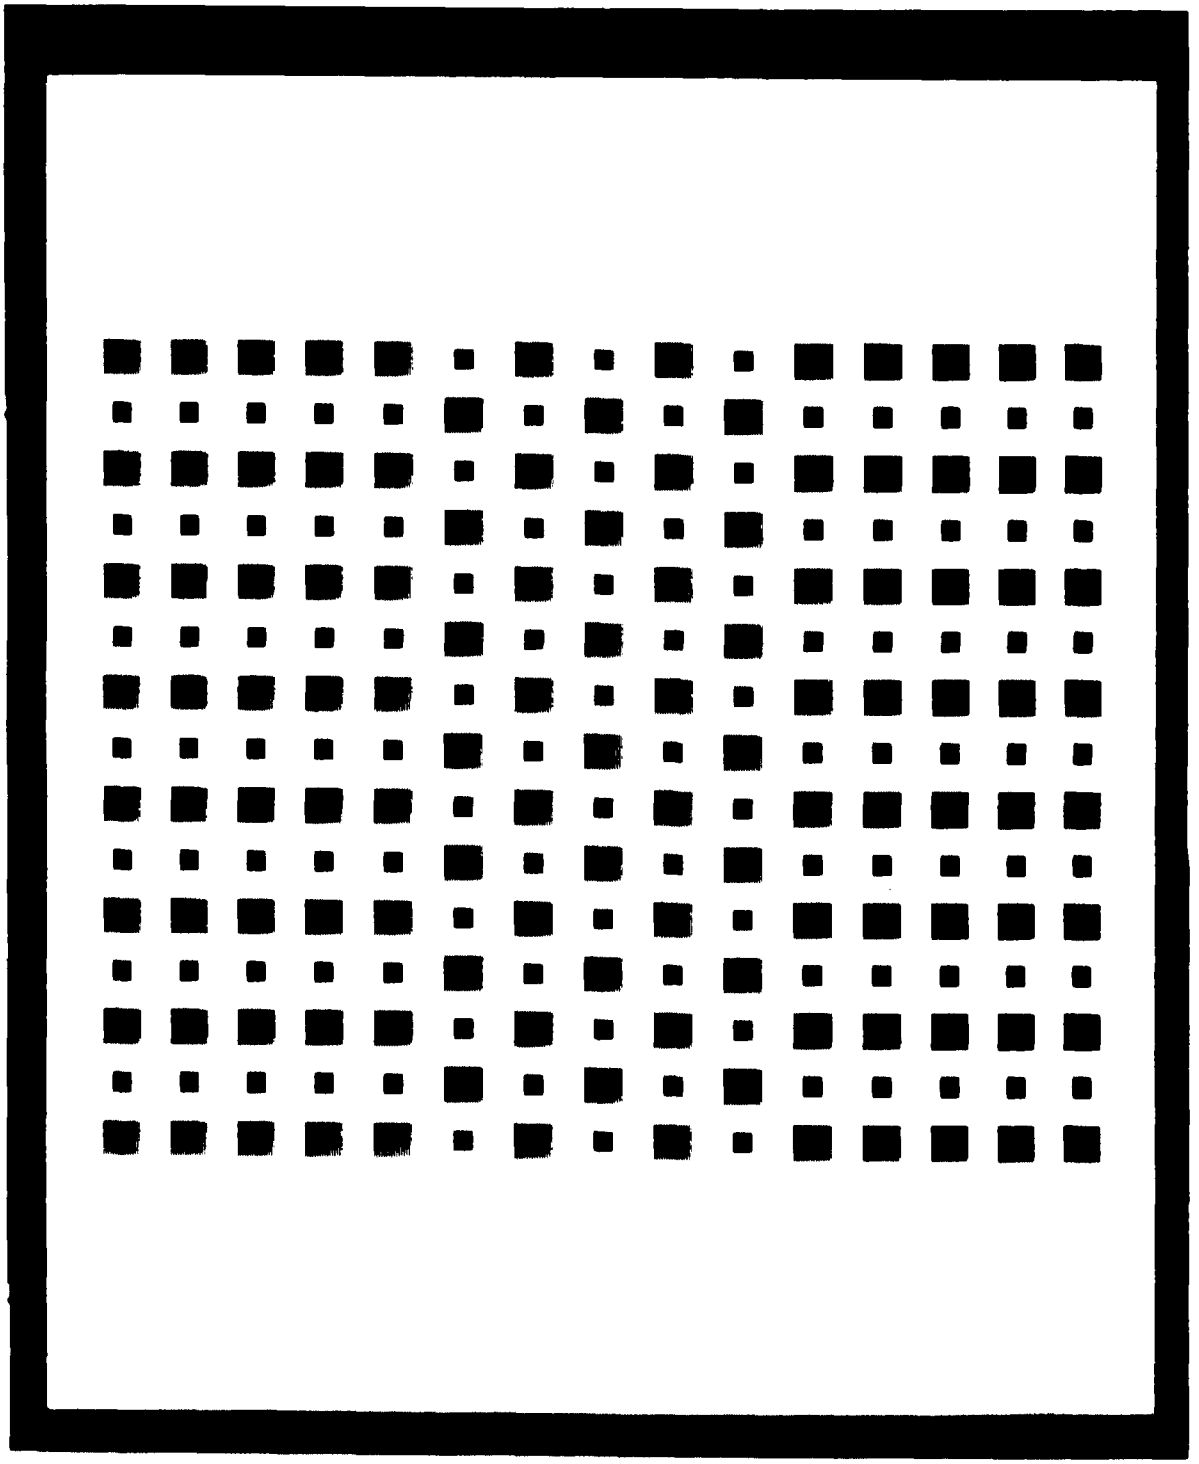

Figure 15

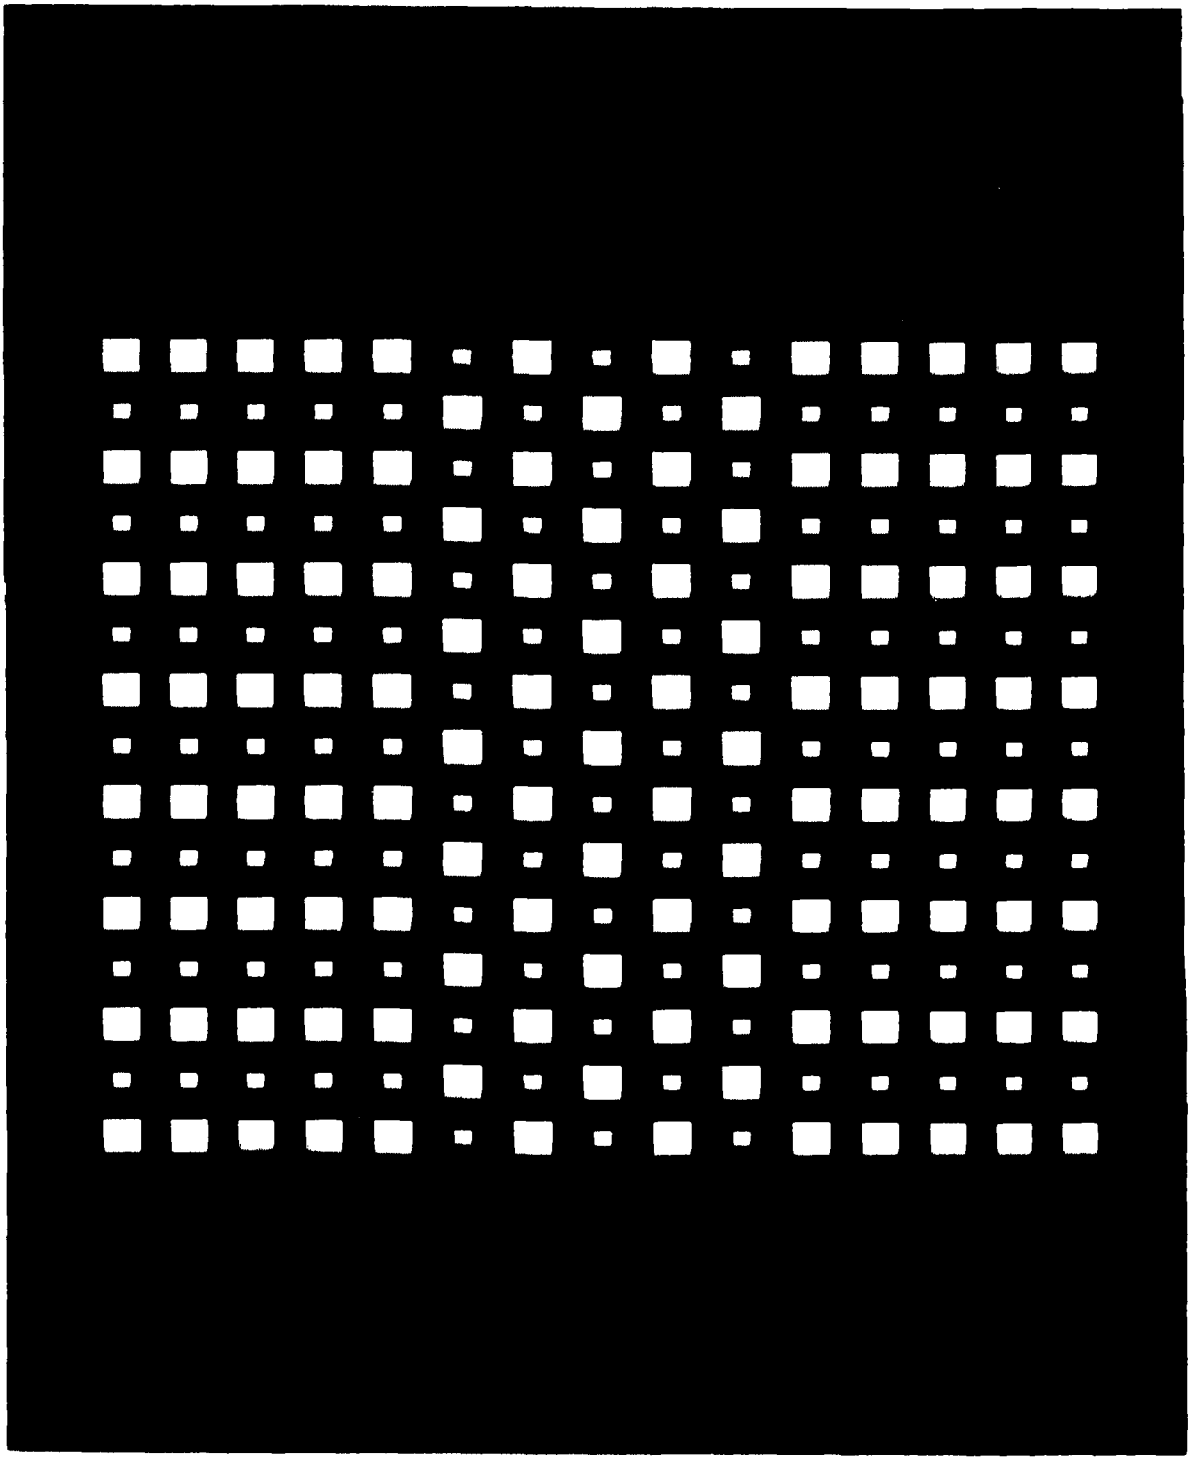

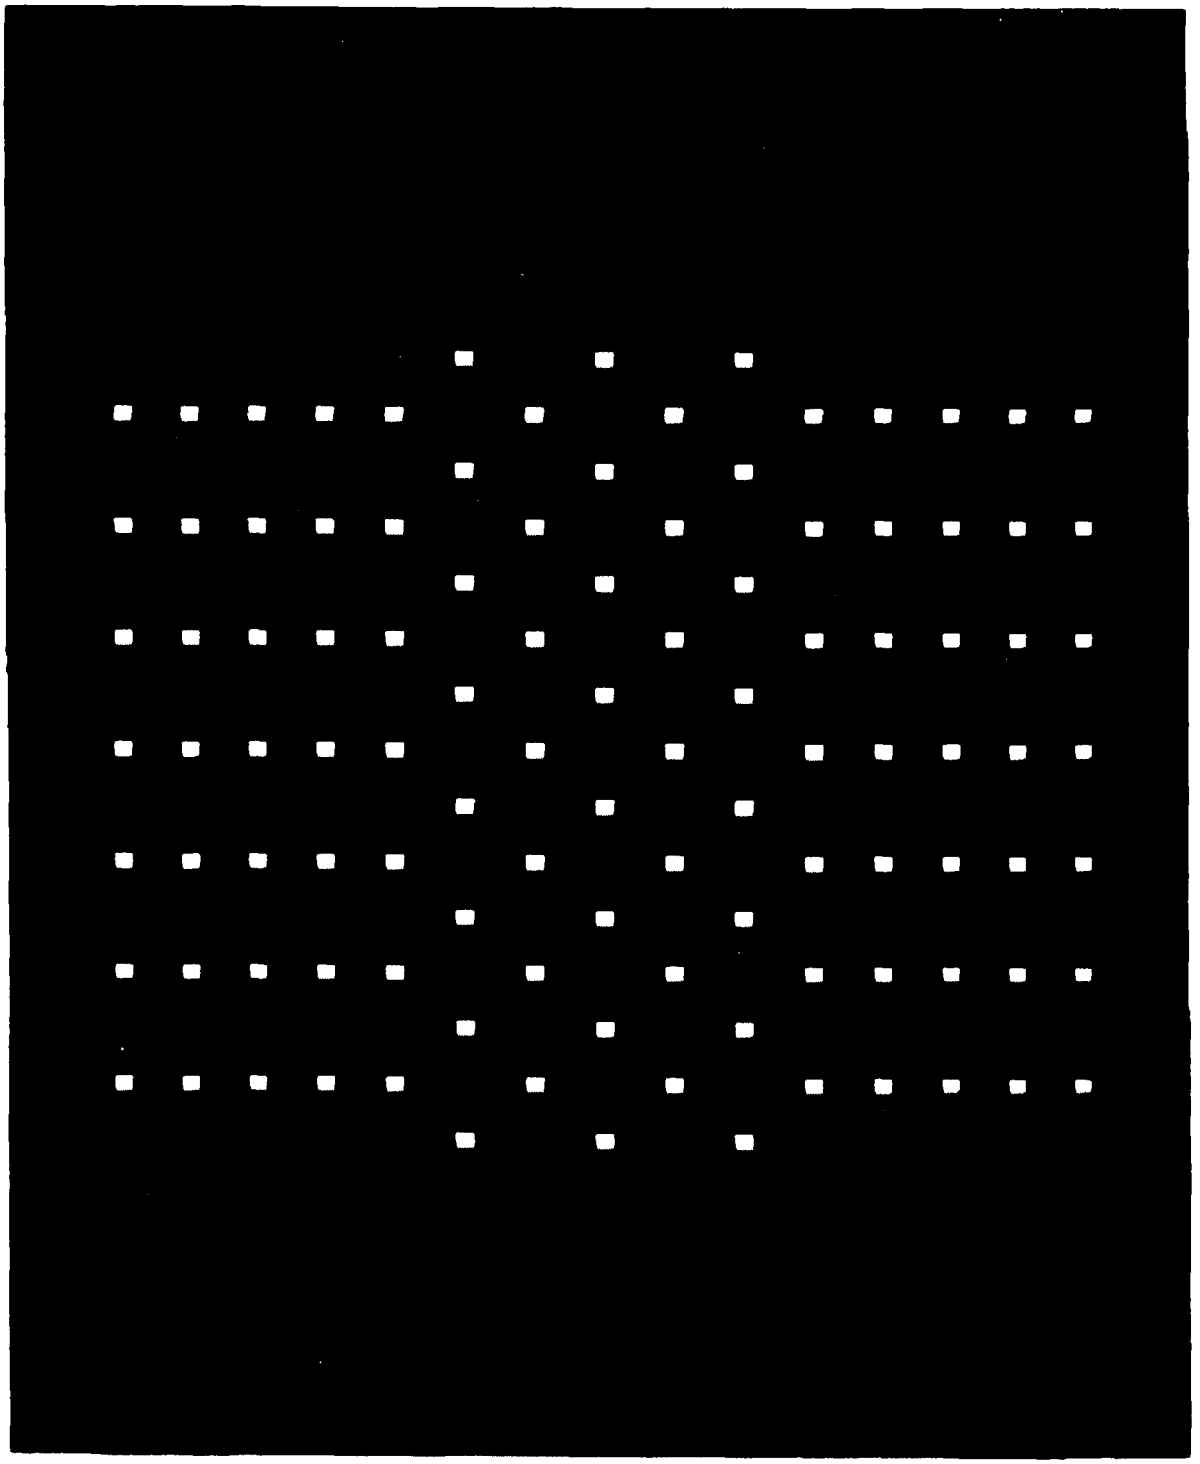

Figure 17

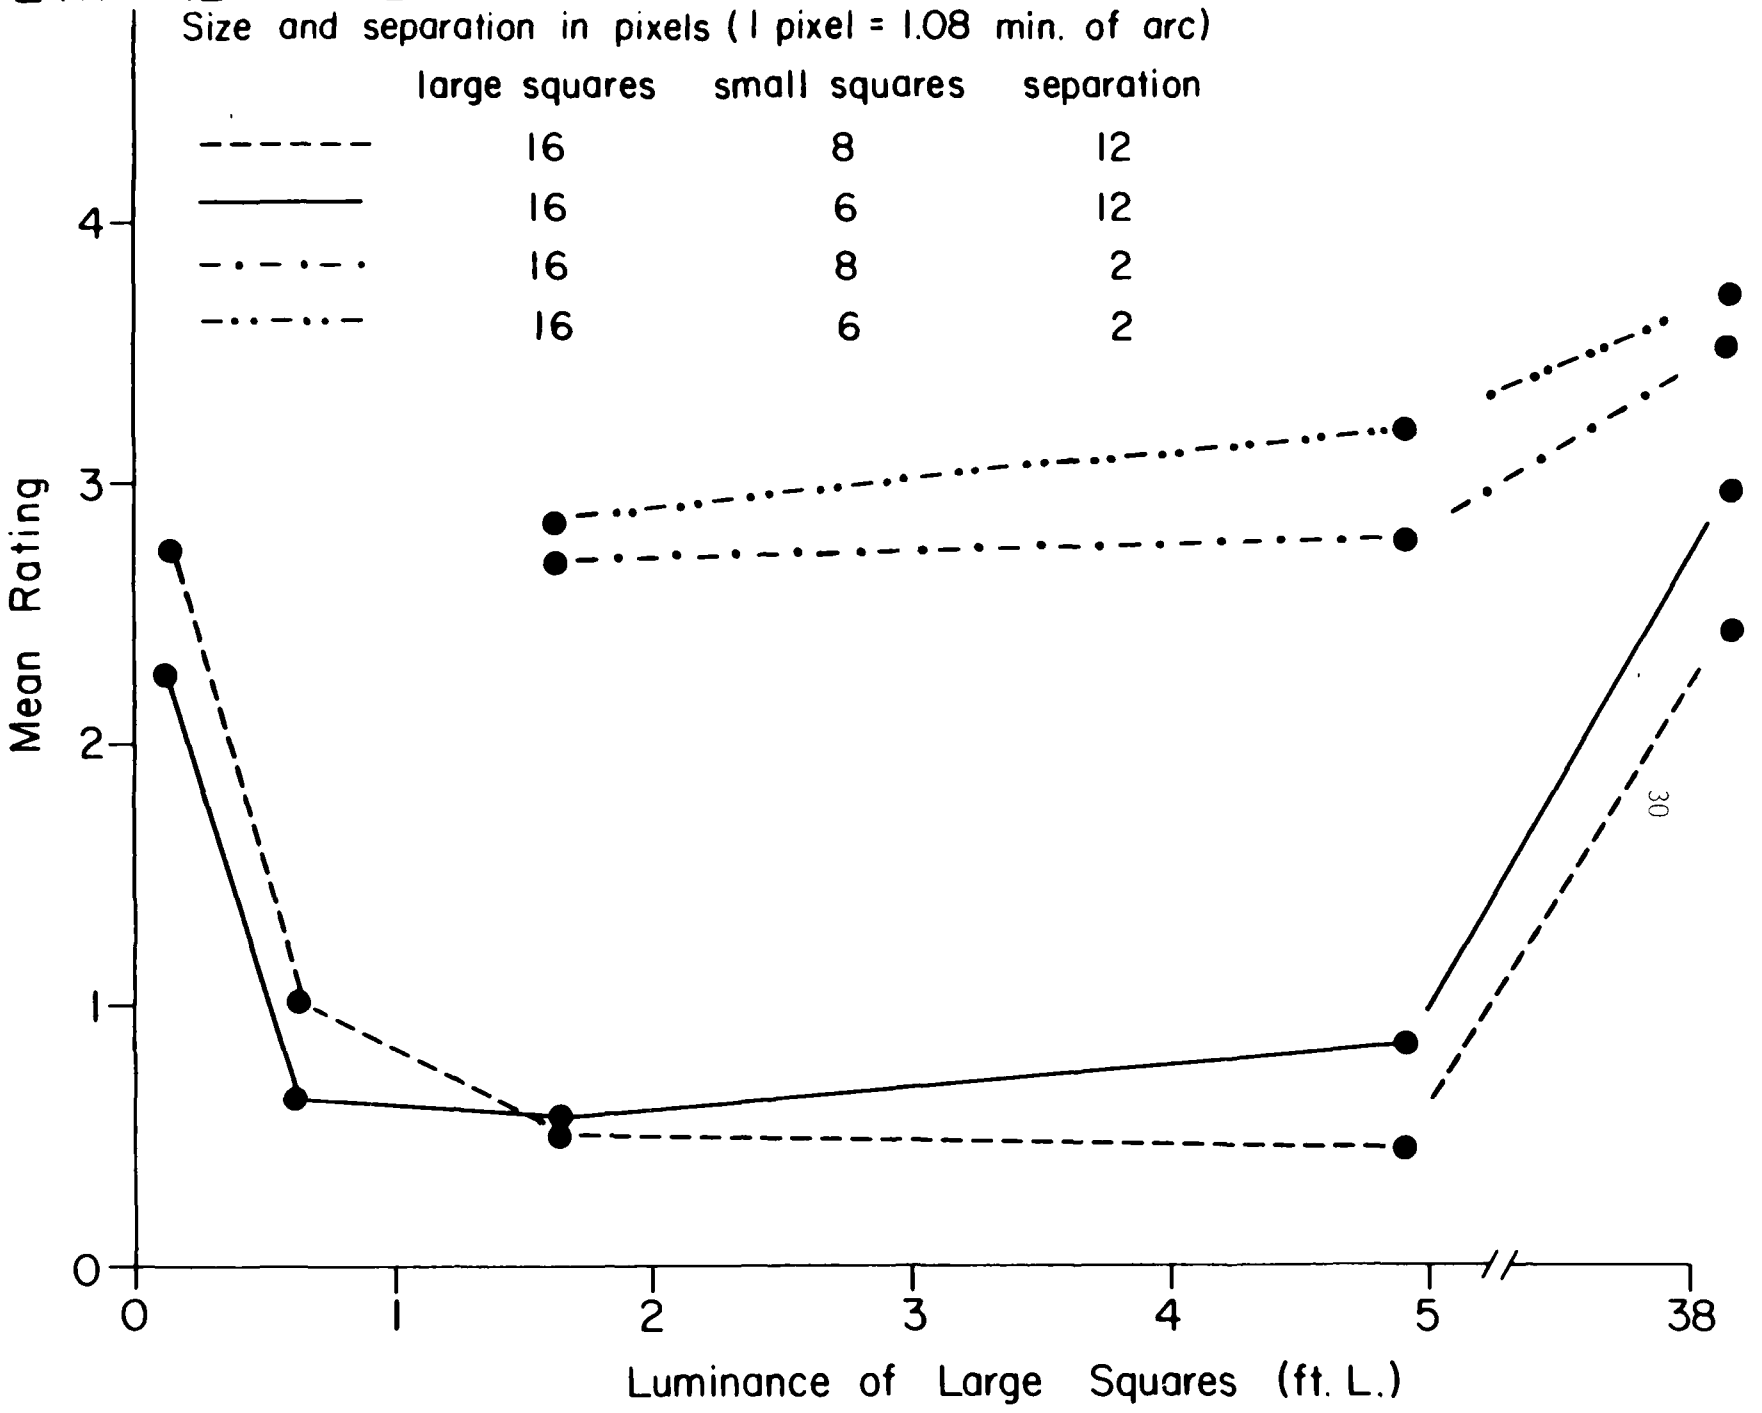
Size and separation in pixels (1 pixel = 1.08 min. of arc)
large squares
small squares
separation
16
8
12
16
6
12
16
8
2
16
6
2
Mean Rating
0
1
2
3
4
0
1
2
3
4
5
38
Luminance of Large Squares (ft. L.)

Figure 18

Increasing the luminance of the large squares improved segregation for both the 8 pixel and 6 pixel small squared displays. Also, as would be expected, tripartite segregation is stronger with the 6 pixel small squares than with the 8 pixel small squares. The analysis of variance with a 12 pixel separation revealed that luminance [$F(4, 36) = 55.5$; $p < .01$] and the luminance by size interaction [$F(4,36) = 24.6$; $p < .02$] were significant. The main variable of size was not significant [$F(1, 9) = .43$]. With a 12 pixel separation, decreasing the luminance of the large squares lessened tripartite segregation for both the 8 pixel and 6 pixel small squares, but not as much for the 6 pixel square as for the 8 pixel square. We should expect a crossover between the mean ratings for the 8 pixel small square pattern and the 6 pixel small square pattern. Above the equal energy point, the energy of the 8 pixel small square is more similar to the energy of the 16 pixel large square, while below the equal energy point, the energy of the 6 pixel small square is more similar to the eneergy of the 16 pixel large square. The ratio of the areas of the 8 and 6 pixel squares is 1.8. The ratio of the large-square luminances giving minimum tripartite segregation for the 8 pixel and 6 pixel small squares was 2.7.

Experiment 8. Experiment 8 was undertaken to examine whether tripartite segregation would occur when low spatial frequencies are removed. The displays consisted of the 16 pixel squares and 8 pixel squares on a black background, .02 ft.-L. The stimuli were: (1) an unfiltered display in which the luminances of the large and small squares were 38.0 ft.-L., (2) this display in which all frequencies below 14 cycles per degree were removed, (3) a display in which the luminances of the small squares were 38.0 ft.-L. and the luminance of the large squares were 14.3 ft.-L., (4) this display in which all frequencies below 14 cycles per degree were eliminated. Seven subjects served in the experiment. Each subject made 10 judgments of each of the displays. In addition to the four stimulus displays four additional displays , two filtered and two unfiltered were added to increase the size of the stimulus set. A stimulus presentation consisted of two filtered and unfiltered displays. A stimulus presentation consisted of a fixation stimulus for 2000 msec., a mask for 250 msec., a stimulus for 1000 msec., and a post-stimulus mask which continued until the subject made their scaling response.

The mean of subjects ratings for the unfiltered equal luminance stimulus was 2.4 with a standard deviation of .54 and when the large squares were made brighter 1.2 with a standard deviation of 1.4. Figure 19 shows a high pass filtered image of the stimulus in which the large and small squares are equal in intensity. All frequencies below 14 cycles per degree have been removed. Tripartite segregation is greatly reduced. The mean of subjects' tripartite segregation ratings was .69 with a standard deviation of .36. Figure 20 shows a highpass filtered image of the stimulus in which the large squares was set at the mean intensity for which tripartite segregation was judged to be minimal in Experiment 6. The mean of subjects segregation ratings are now 3.3 with a standard deviation of .66. Thus, if a pattern which fails to produce segregation because the large and small squares have been adjusted to the same average contrast is filtered to remove all spatial frequencies below 14 c/deg, the resulting pattern is immediately seen as segregated. High spatial-frequency information in the absence of conflicting low spatial-

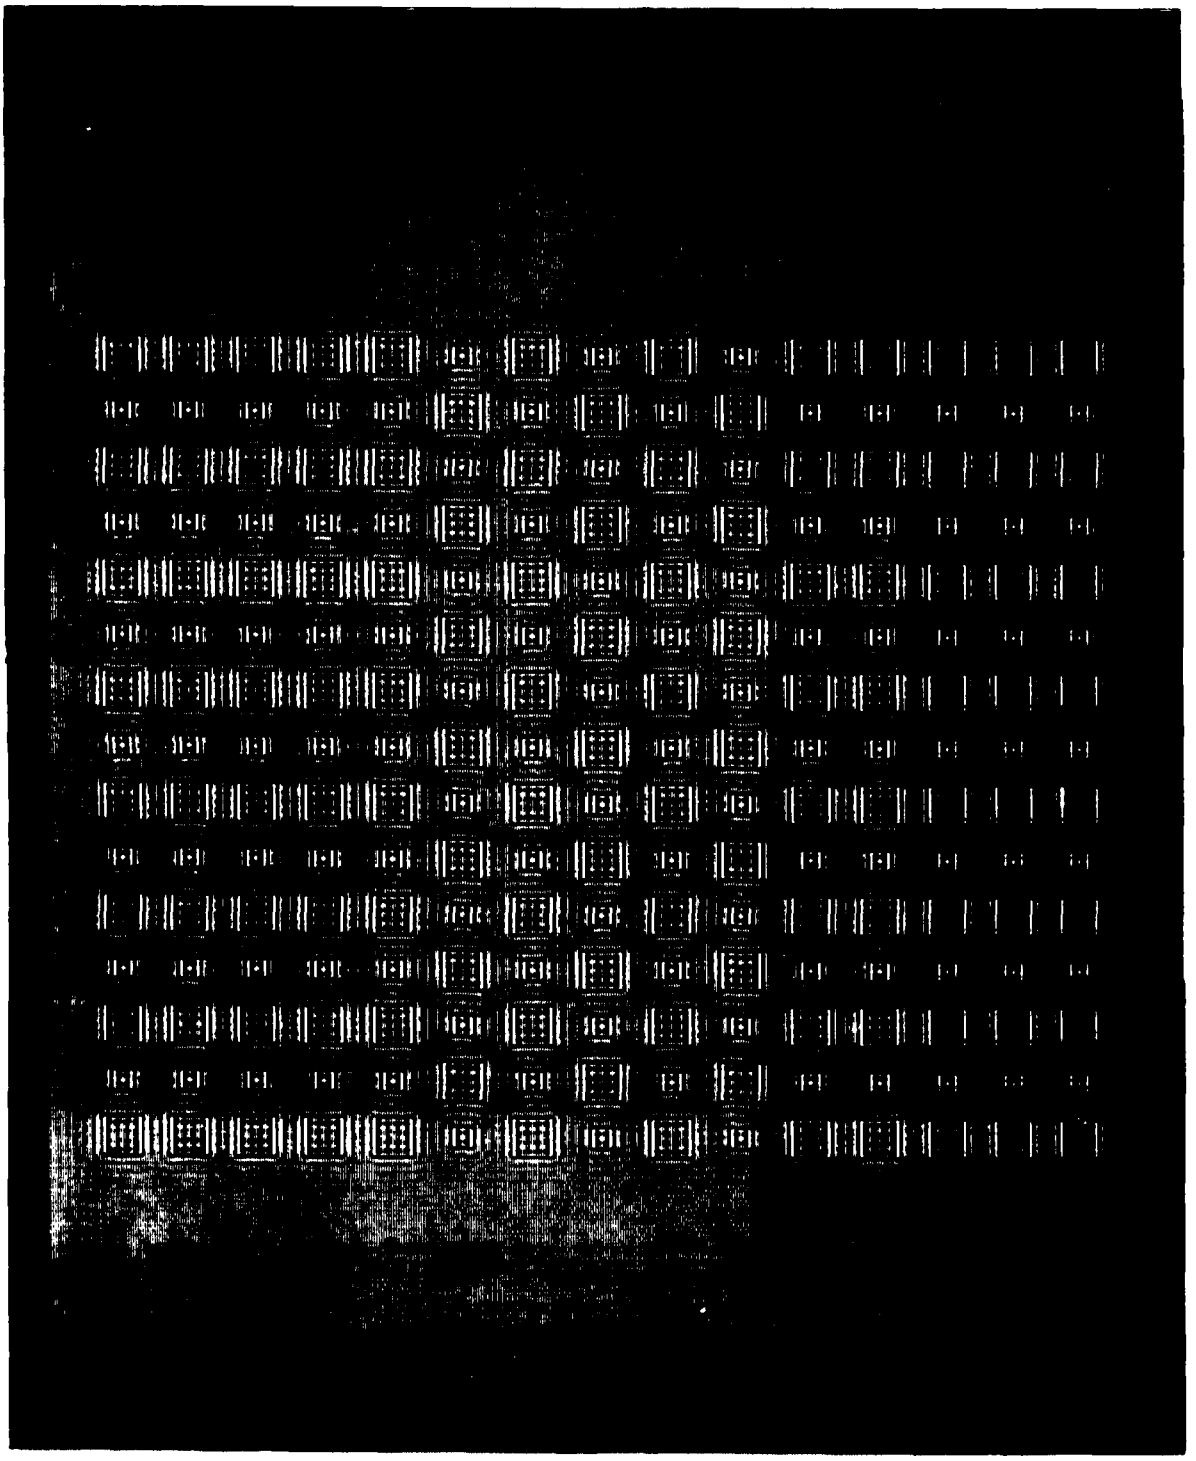

Figure 19

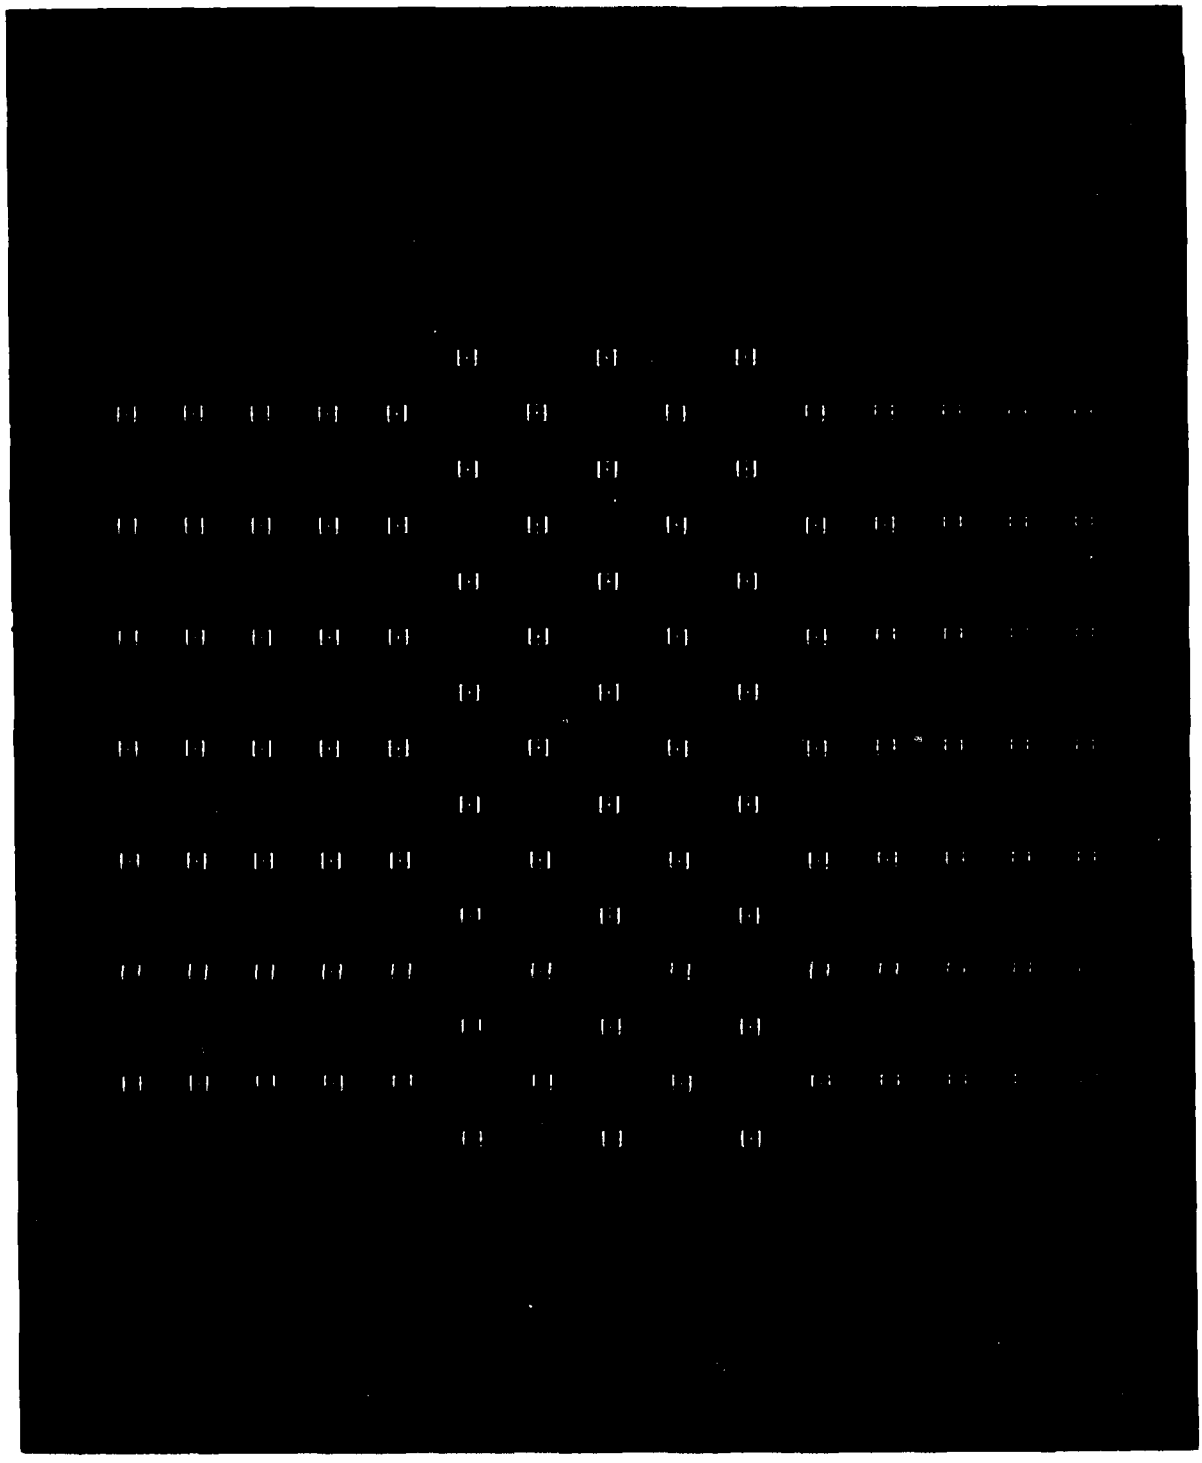

frequency information may be sufficient to produce tripartite segregation even though low-spatial-frequency information - when present - dominates and prevent it.

2.4 Discussion

How are the effects of energy variables to be explained? The results of the experiments with large and small squares shows that a difference in size can be canceled by a difference in contrast. The large dim and the small bright squares would elicit the same magnitude of response from large receptive fields. An observer does not see *tripartite* segregation if the squares' luminances have been adjusted so that their 'energies' (contrast times size) are equal. Since the small and large squares differ in their higher frequencies, this is an example of low frequencies masking both a size difference and higher frequency energy differences. The experimental results with the tripartite display indicate that if there is low spatial-frequency information present in the stimulus, an observer perceives good tripartite segregation if and only if the two squares elicit sufficiently different responses from the low spatial-frequency mechanisms. Bandpass filtering various patterns that do and do not produce tripartite segregation support our hypothesis with the unequal square displays.

Figure 16 shows the large and small squares with the same intensity on a black background. Tripartite segregation occurs. Figure 21 shows the bandpass filtered image when the stimulus is convolved with a difference of Gaussians. The excitatory and inhibiting signals differ by a factor of 2. The excitatory sigma, approximately 4 pixels, was half the width of the small squares, and the inhibitory sigma the width of the small squares. The bandpass filtered image shows how elliptically oriented receptive fields would be stimulated by a pattern. The bandpass filtered image shows, as one would expect, that the top and bottom regions would stimulate vertically oriented receptive fields more strongly than the middle region.

Figure 17 shows the display in which tripartite segregation was reduced by decreasing the intensity of the large squares. Figure 22 shows the result of convolving a stimulus using the ratio of the logarithms of intensities of the large and small squares. The chosen intensity values were equal to the intensity values at which subjects reported tripartite segregation to be minimal. The bandpass filtered image shows that when the intensity of the small squares is increased relative to the intensity of the large squares, elliptical oriented receptive fields in the top and bottom regions and in the middle region would be more equally stimulated. Tripartite segregation is greatly reduced. Figure 23 shows the bandpass filtered image when the separation between the squares is 2 pixels. Since the large squares and the small squares are closer to each other in the vertical direction than in the horizontal direction, spatial averaging would cause vertically oriented bar detectors to be stimulated in the top and bottom regions but not in the middle region. Tripartite segregation no longer depends on the intensity of the background because the visual system signals the presence of vertically oriented blobs in the top and bottom regions but not in the middle region. Oriented blobs have been shown to be important in texture segregation but have never been defined precisely. Low frequency blurring may, in

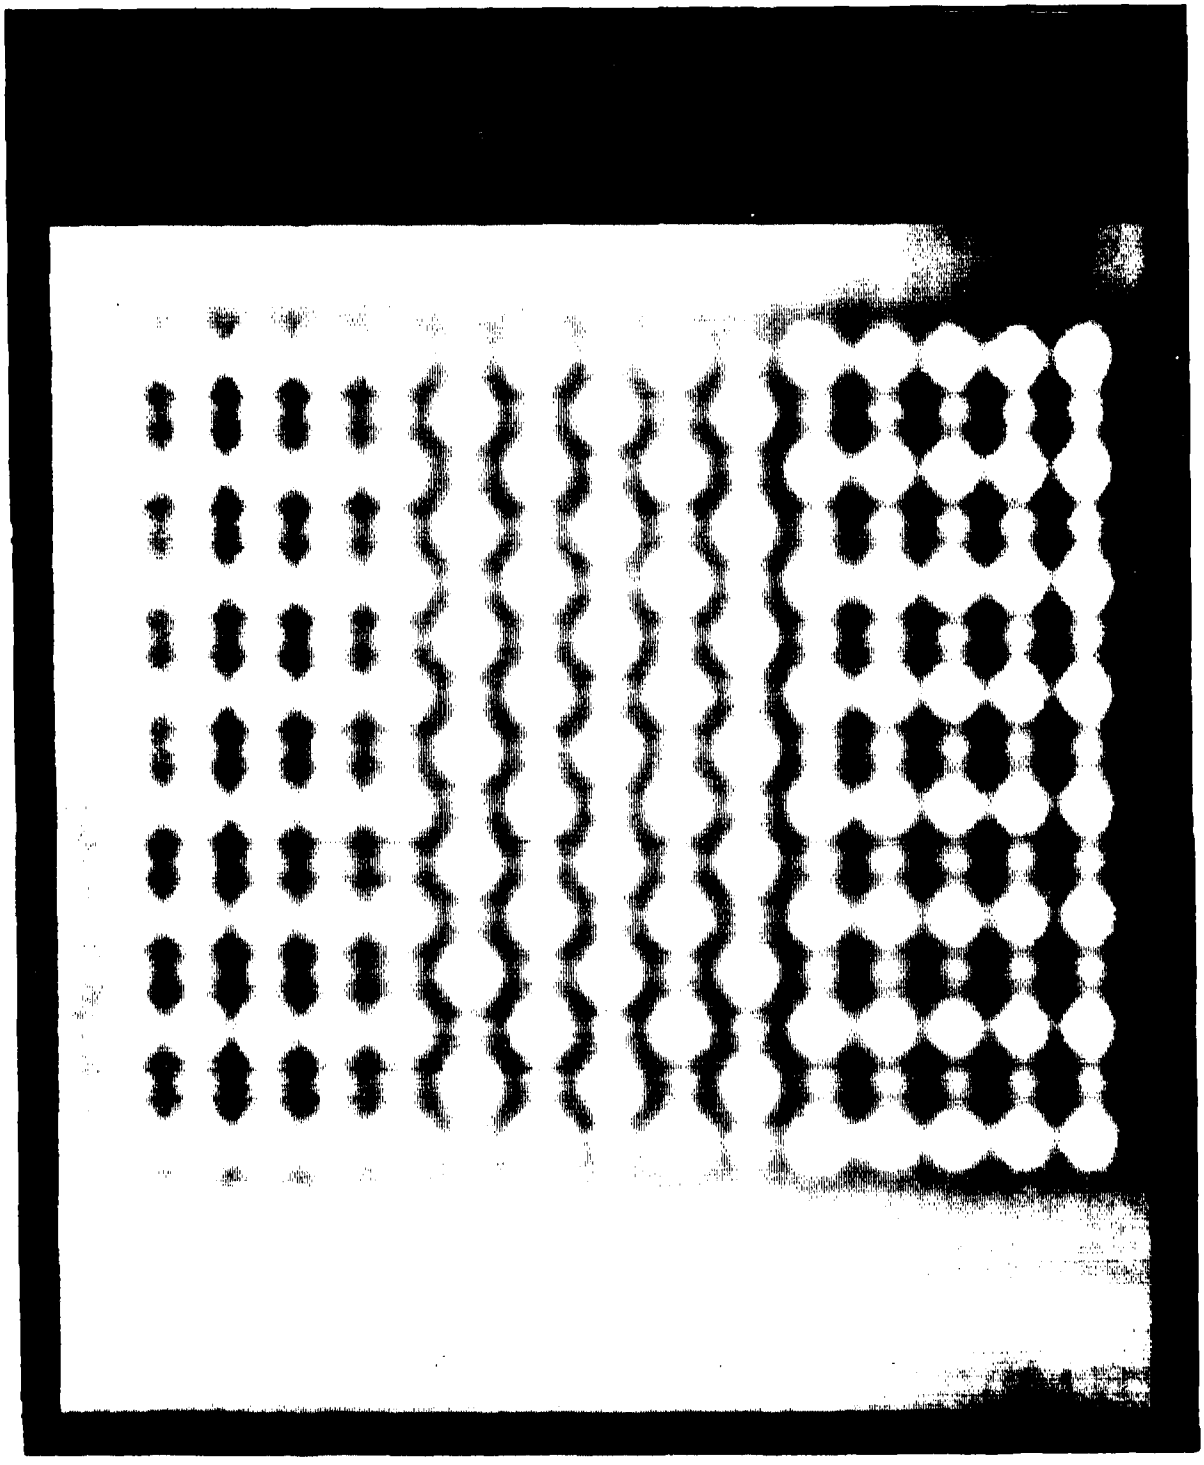

Figure 21

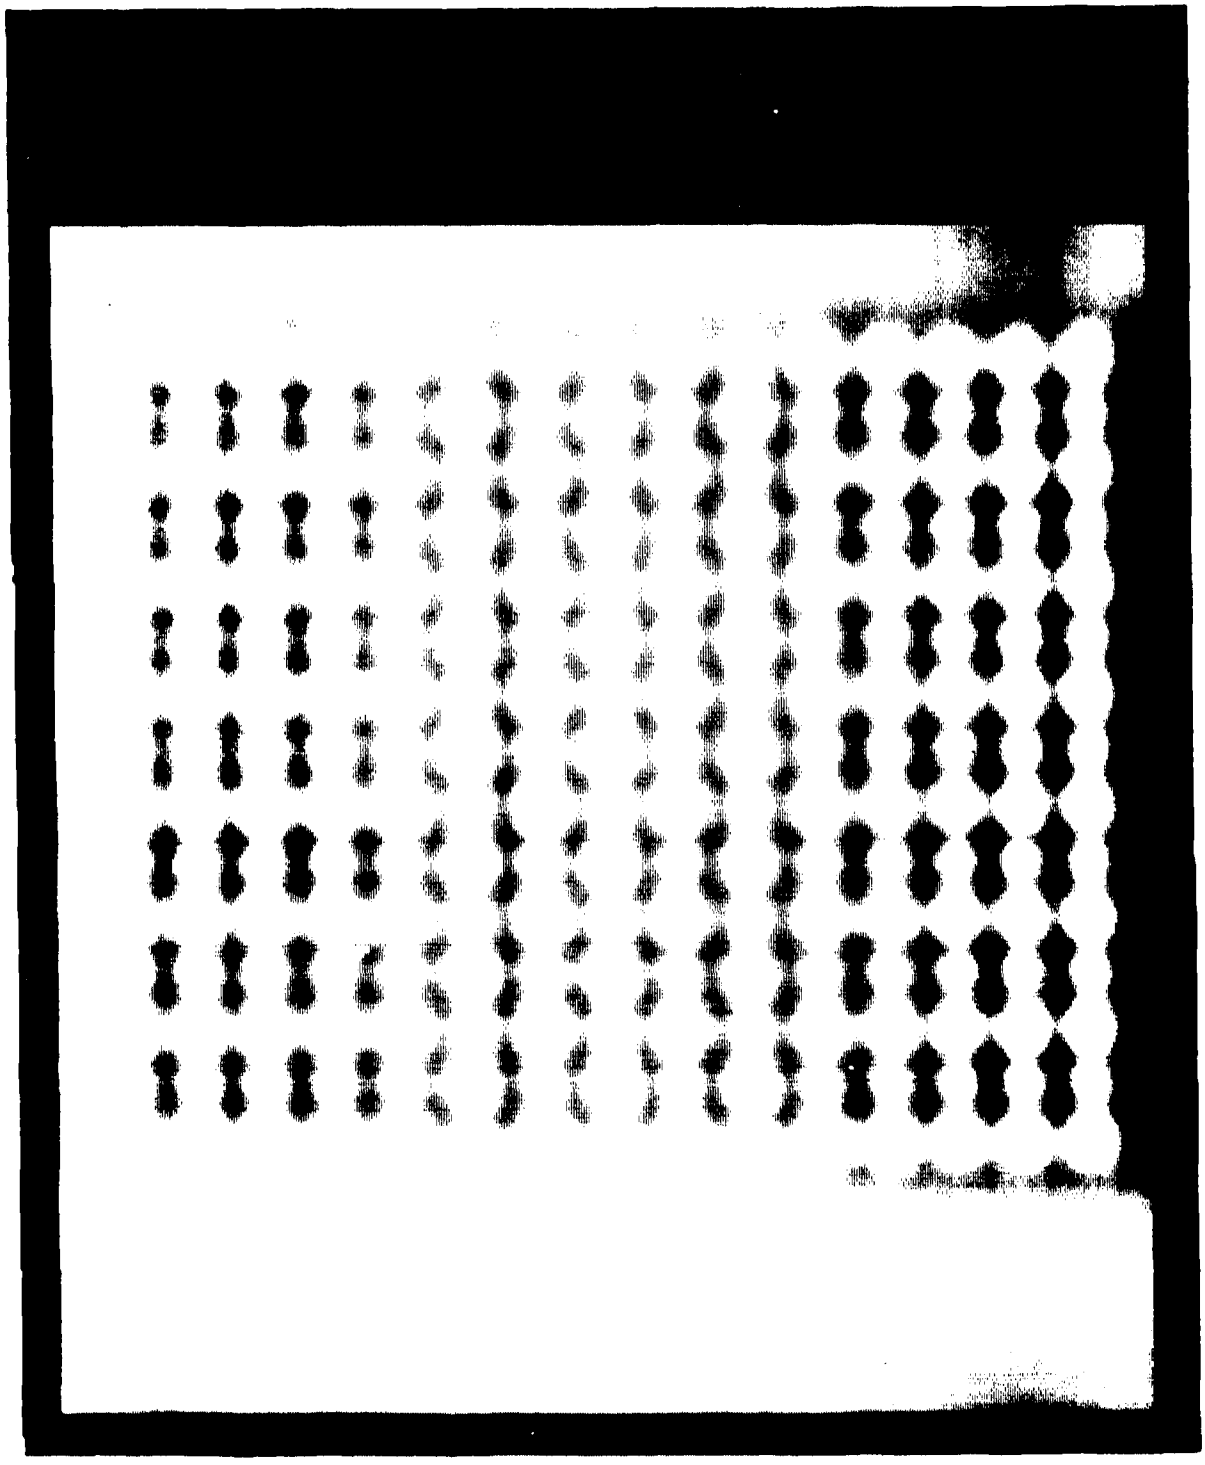

Figure 22

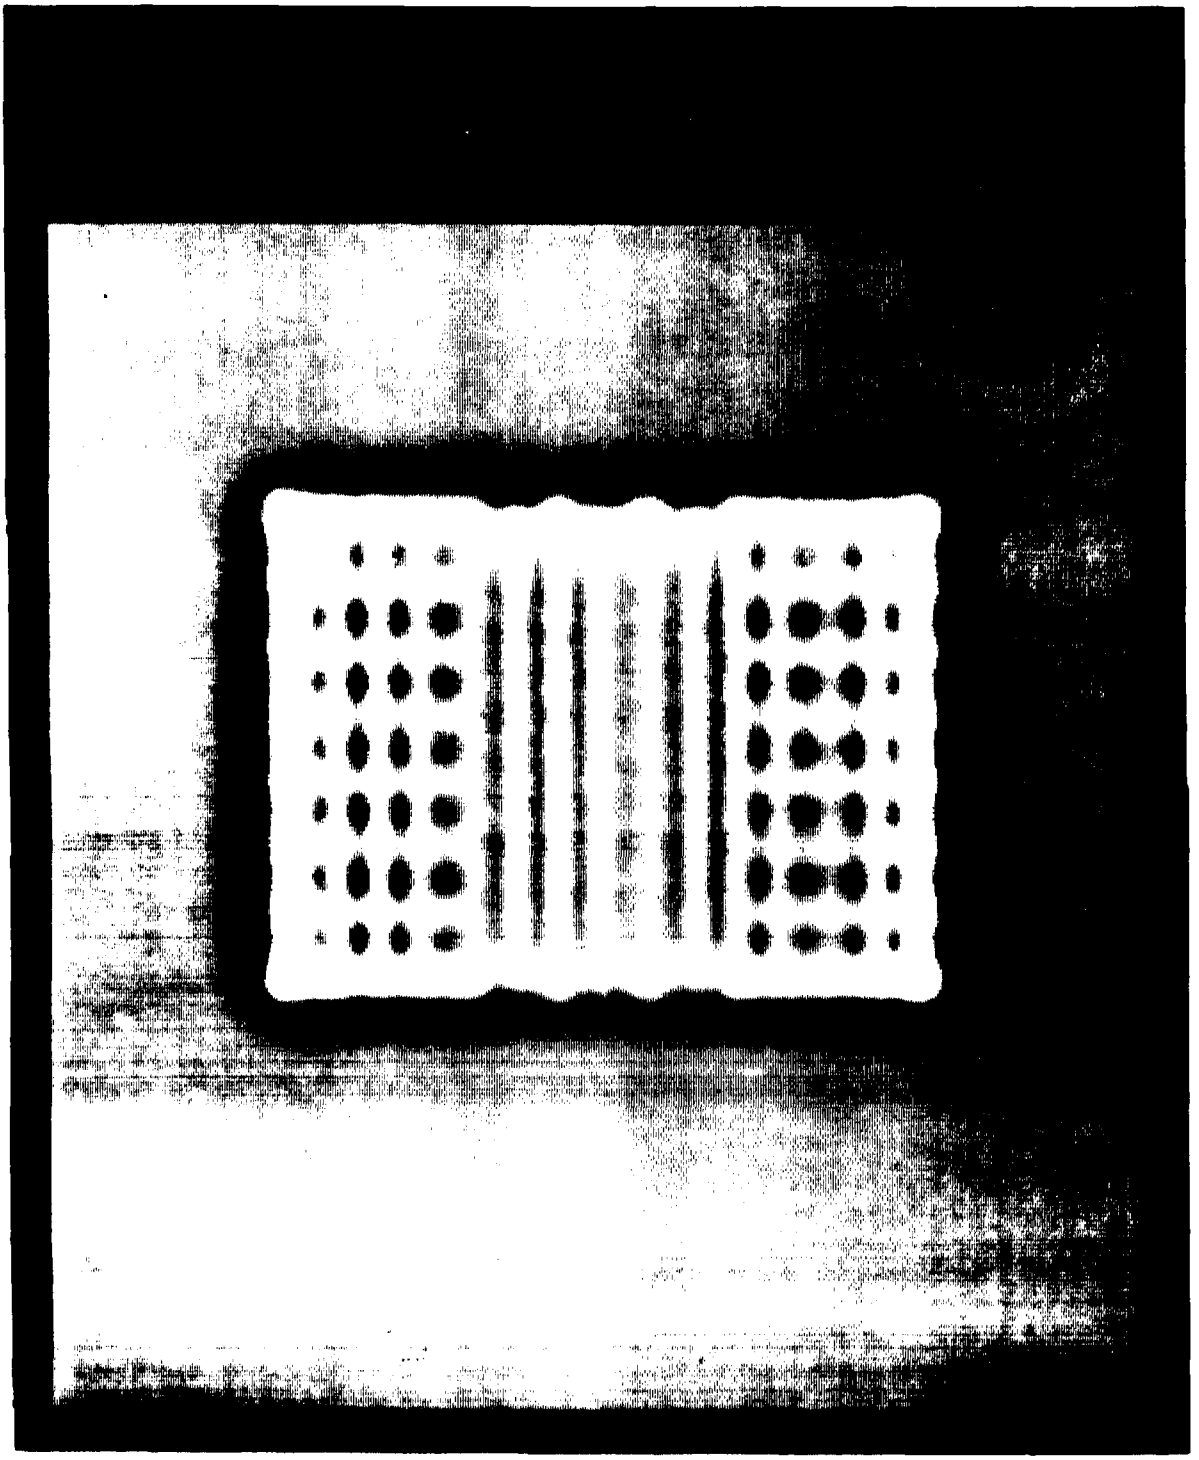

fact, constitute the definition of an oriented blob.

Figure 14 shows the large and small squares with the same intensity on a white background produces tripartite segregation. Figure 24 shows the bandpass filtered image. The image shows that the top and bottom regions would stimulate vertically oriented receptive fields more strongly than the middle region. Figure 15 shows the display in which tripartite segregation was reduced by decreasing the intensity of the small squares. Figure 25 shows the result of convolving a stimulus using the ratio of the logarithms of intensities of the large and small squares. The chosen intensity values were equal to the mean values at which subjects reported tripartite segregation to be minimal. The bandpass filtered image shows that when the intensity of the small squares is decreased relative to the intensity of the large squares, elliptical oriented receptive fields in the top and bottom regions and the middle region would be more equally stimulated. Tripartite segregation is greatly reduced.

Figure 26 shows a tripartite pattern in which the small squares are 6 pixels rather than 8 pixels. Figure 27 shows the same pattern in which the large squares have been reduced to the same intensity as in the previous pattern in which the small squares were 8 pixels. Tripartite segregation is stronger with the 6 pixel small squares than with the 8 pixel small squares (compare Figures 27 and 17). Figure 28 shows the bandpass filtered image. As one would expect, vertically oriented receptive fields are stimulated more strongly in the top and bottom regions than in the middle region. Figure 29 shows the tripartite pattern when the intensities of the large squares have been further reduced. Tripartite segregation is lessened. Figure 30 shows the bandpass filtered image. When elliptically oriented receptive fields are stimulated equally in the top, middle and bottom regions, tripartite segregation is minimal.

We now turn to the experiments with equal size squares. Figure 31 shows that when the background intensity is between tripartite segregation occurs strongly even when the high- to low-square intensity ratio is small. Sign of contrast is a feature and averaging occurs separately for positive and negative contrasts. Figure 2 shows that differences in lightness which are easily seen when the lightnesses of individual squares are compared fail to give tripartite segregation when the contrast sign is the same. With contrast sign the same, tripartite segregation occurs only when the ratio of the background intensity to the high-square intensity is less than 1.7 (see Figure 7). Large contrast differences are also required except when the contrast difference between the background and the high-square is small. In that case, bandpass filtering destroys the lighter square. Averaging causes the lighter squares to be absorbed by the background. Magnitude of contrast is not a feature. Tripartite segregation occurs only when there is a difference in the distributions of the different size receptive fields stimulated in the top and bottom and in the middle regions.

When the background intensity was above, tripartite segregation occurred strongly only with a 2 pixel separation when the high- to low- square ratios were 1.4 and 1.8 (see Figure 11). Rated segregation decreased steeply with larger separation. Figure 9 shows a stimulus with a 2 pixel separation on a white background, and Figure 32 shows the bandpass filtered image. The

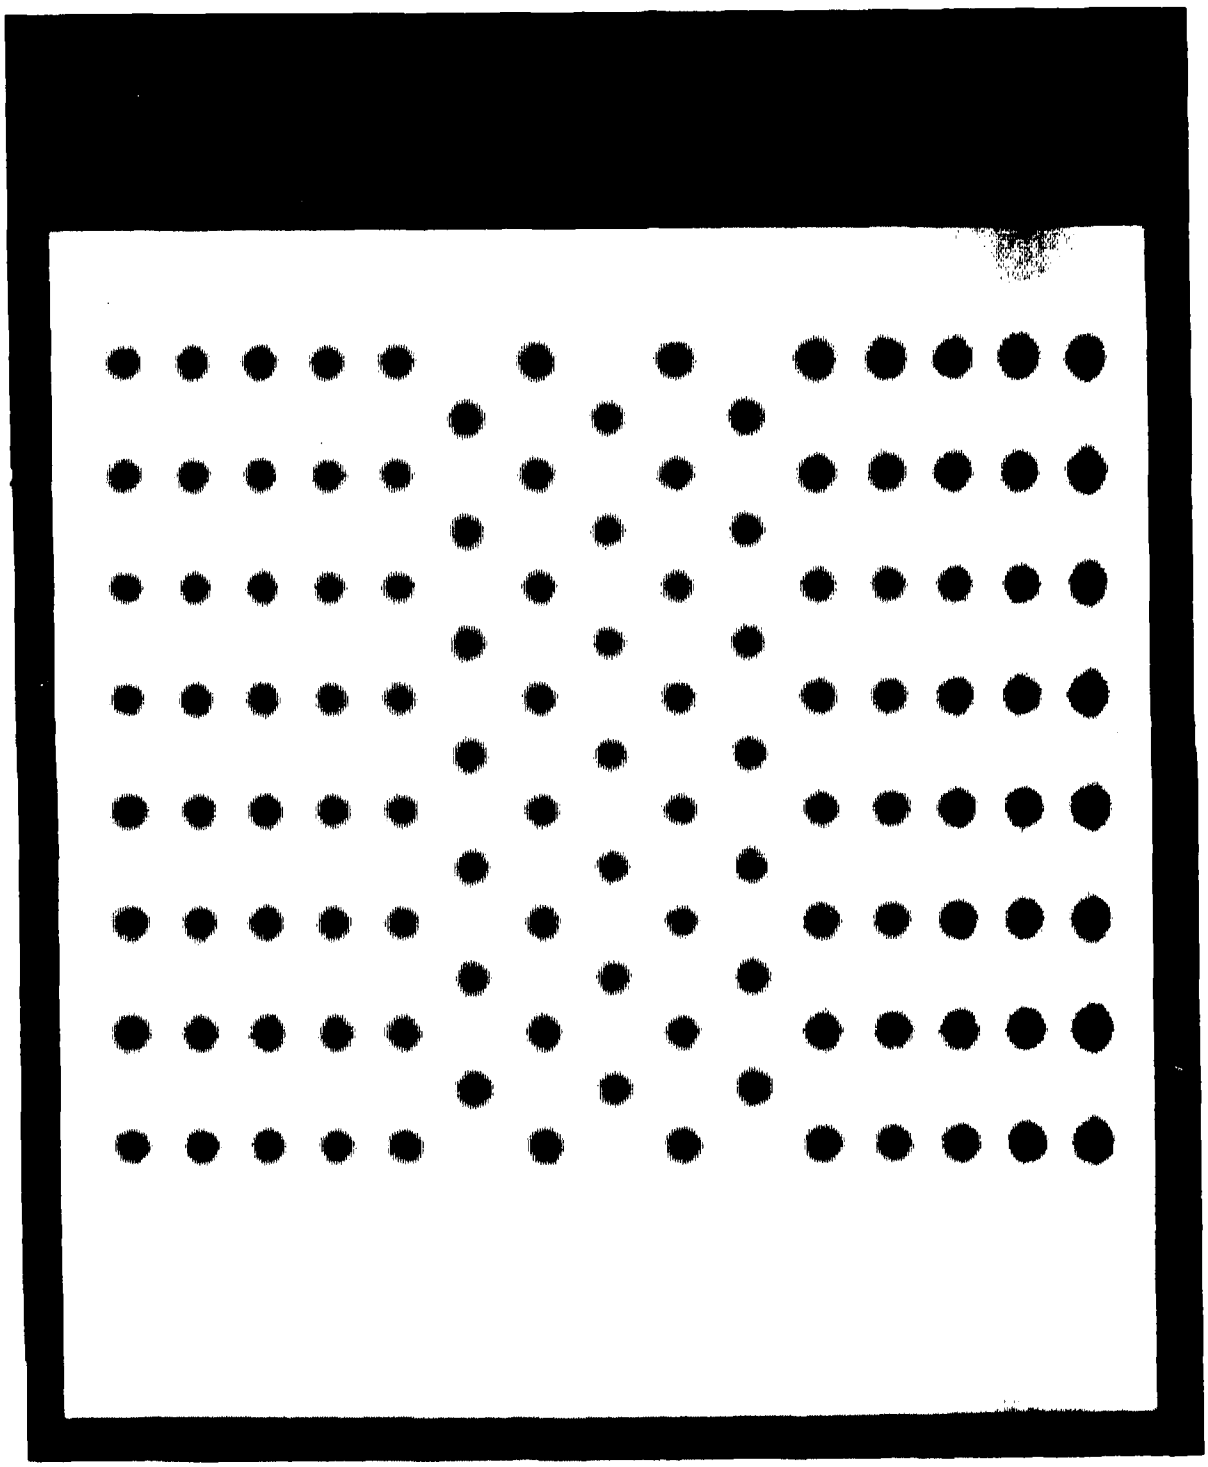

Figure 24

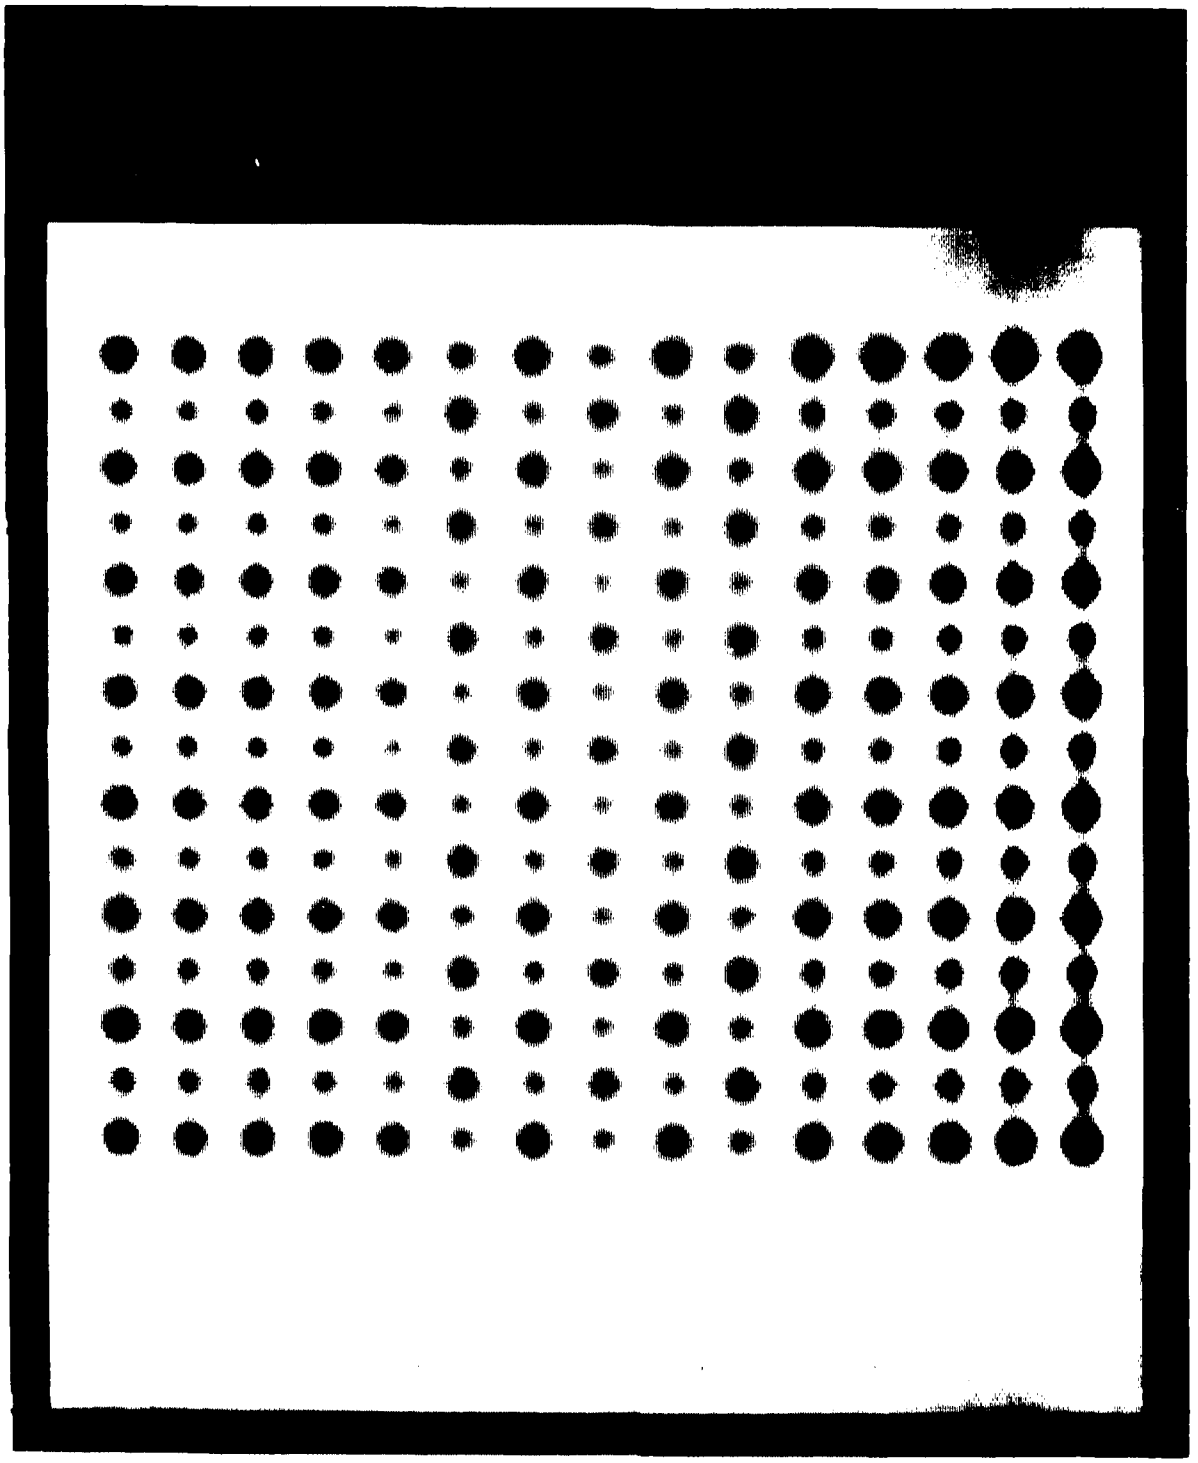

Figure 25

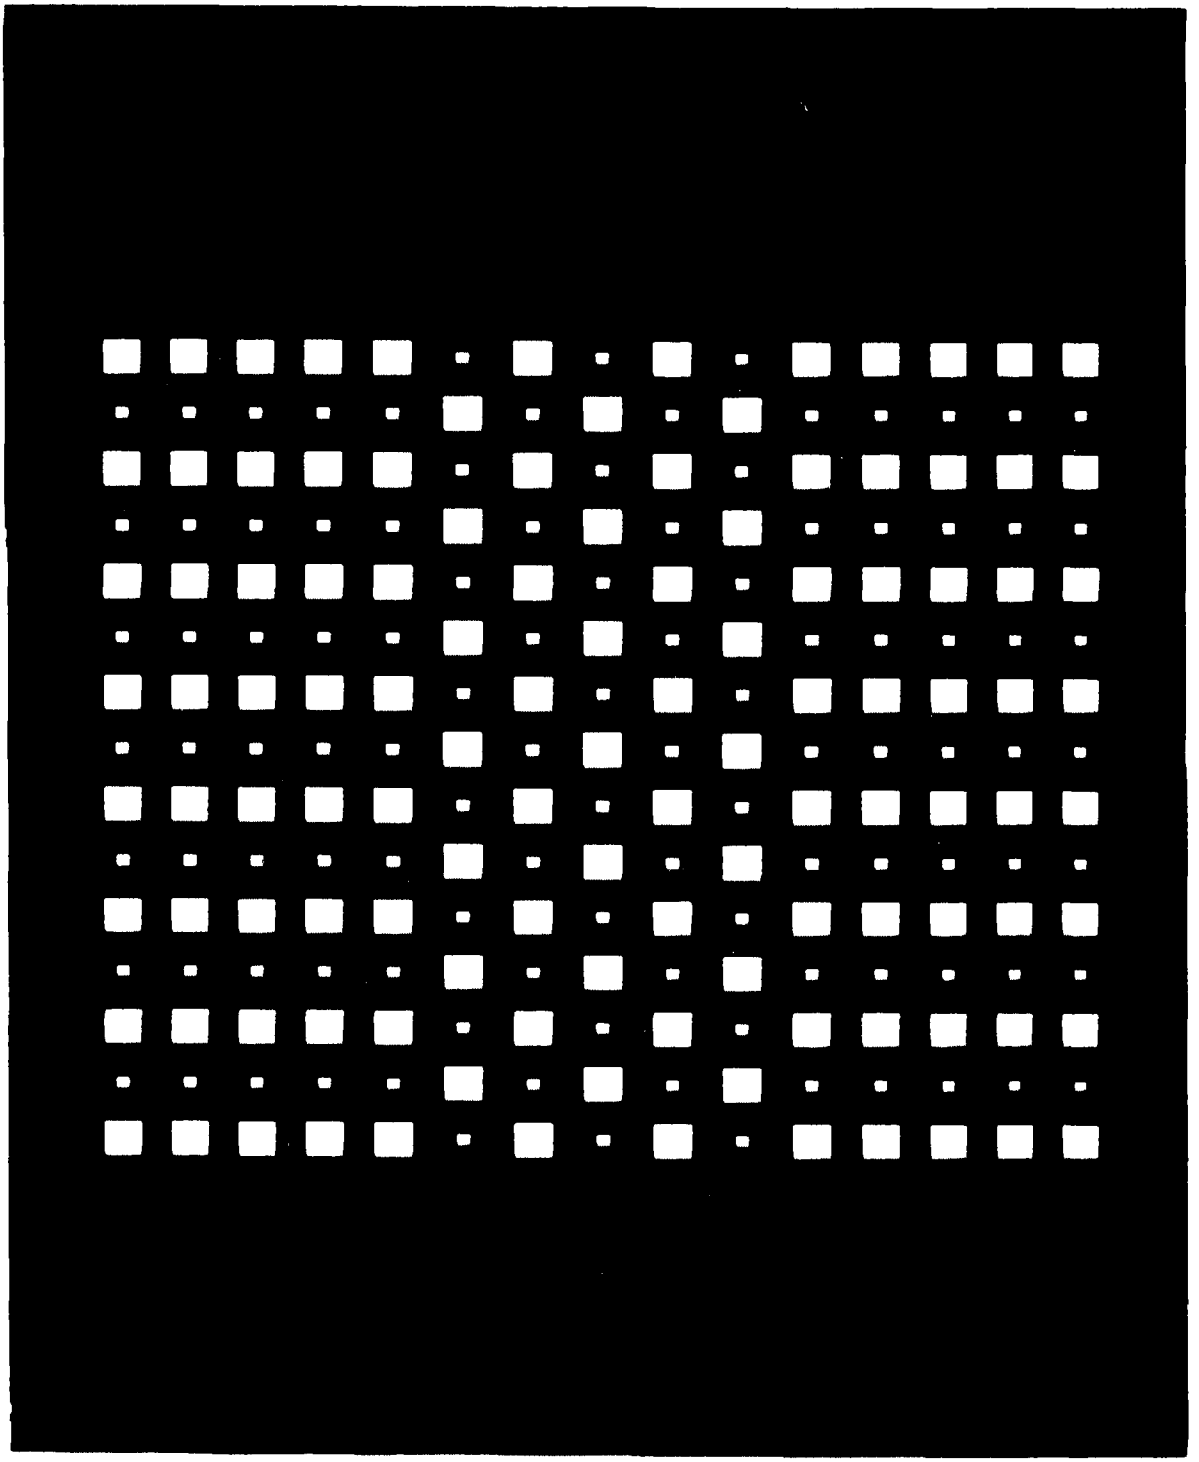

Figure 26

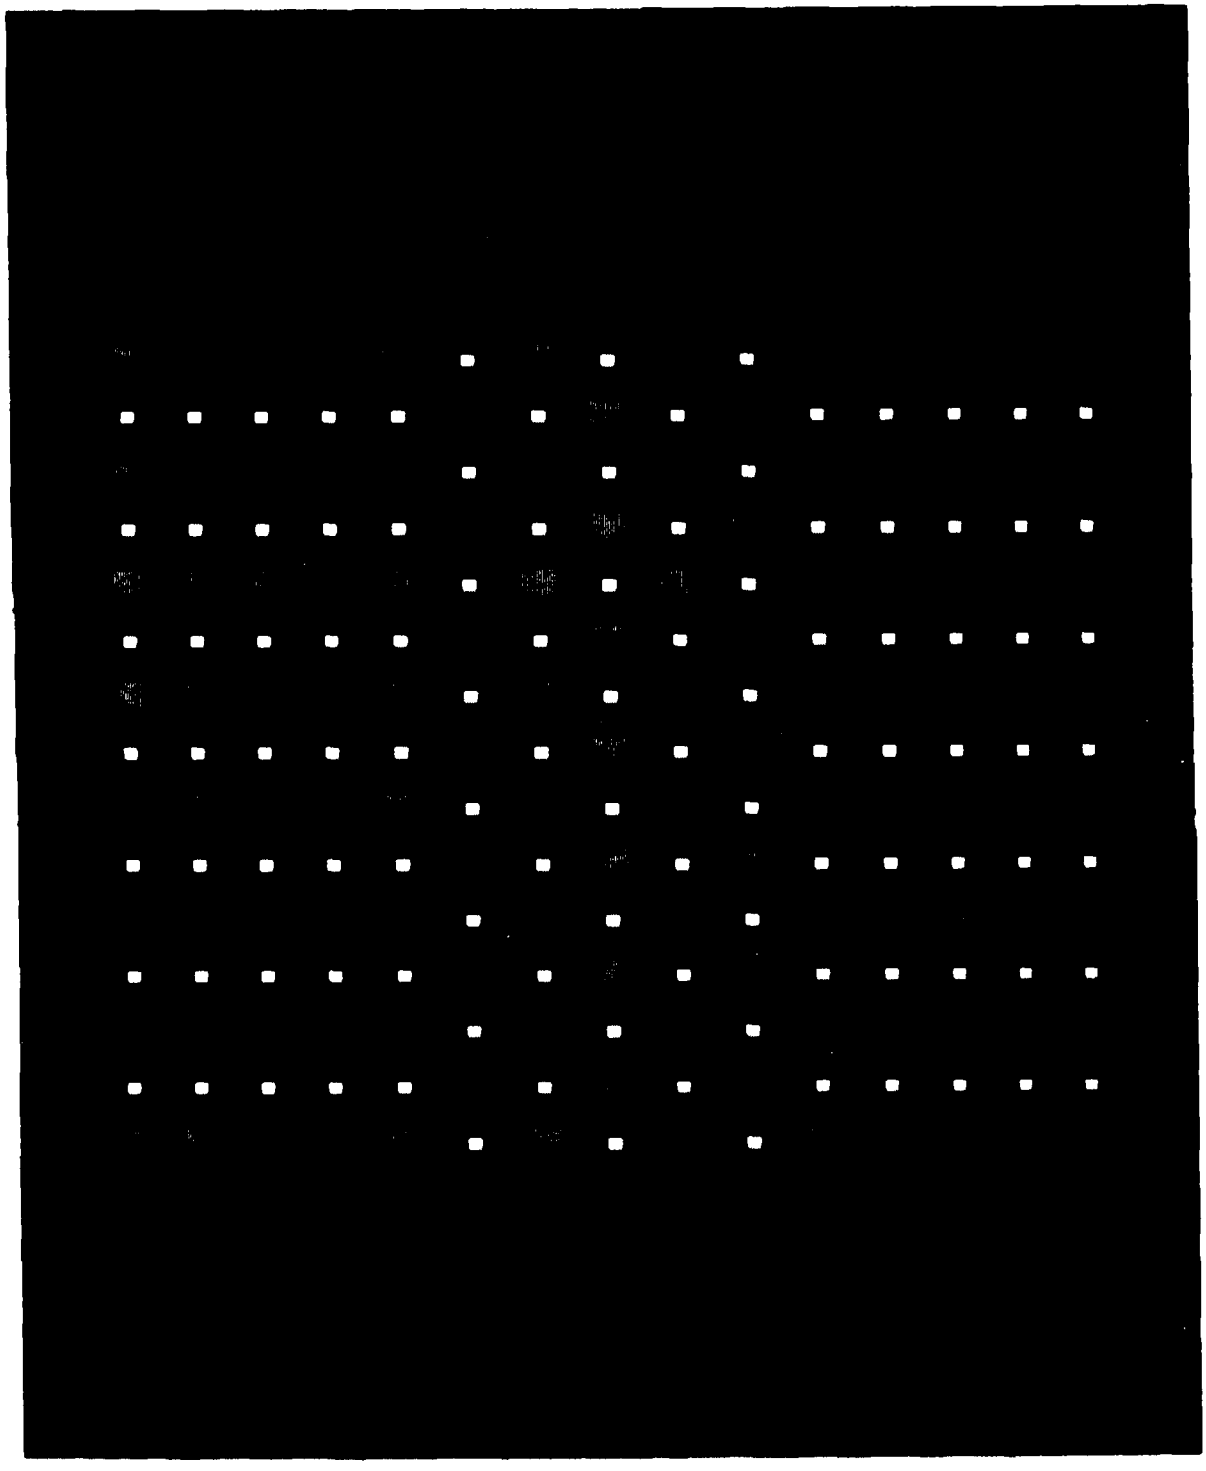

Figure 27

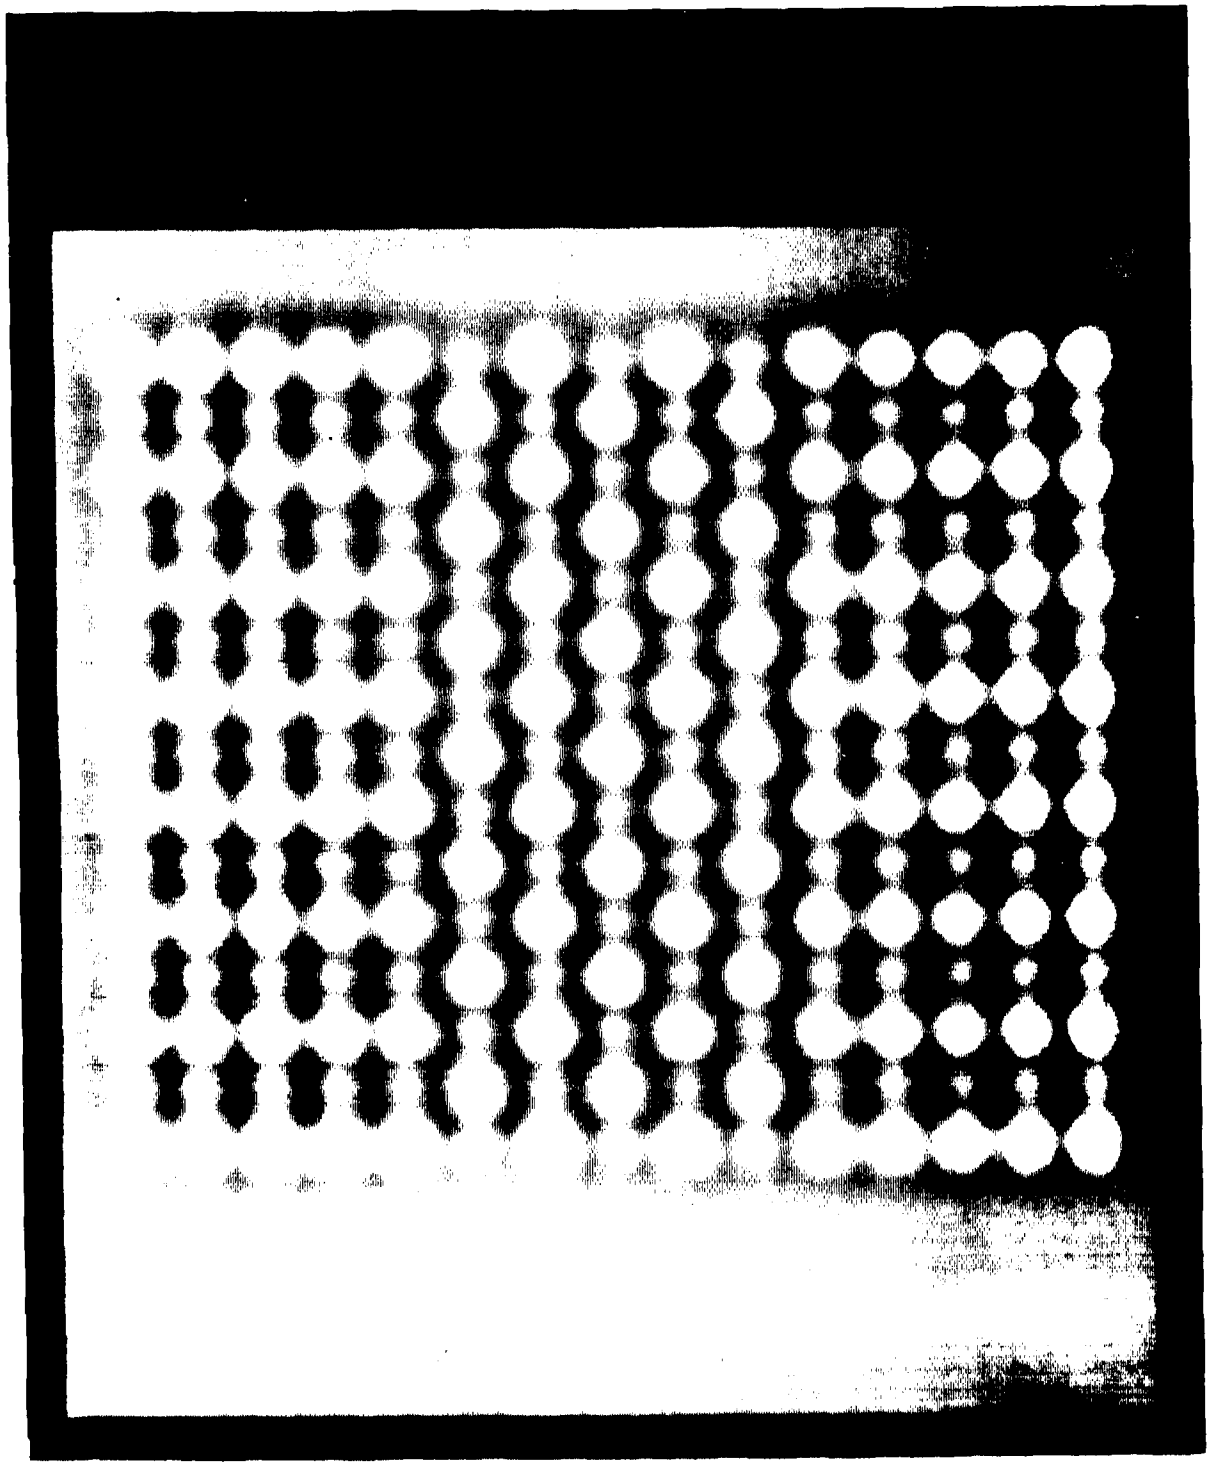

Figure 28

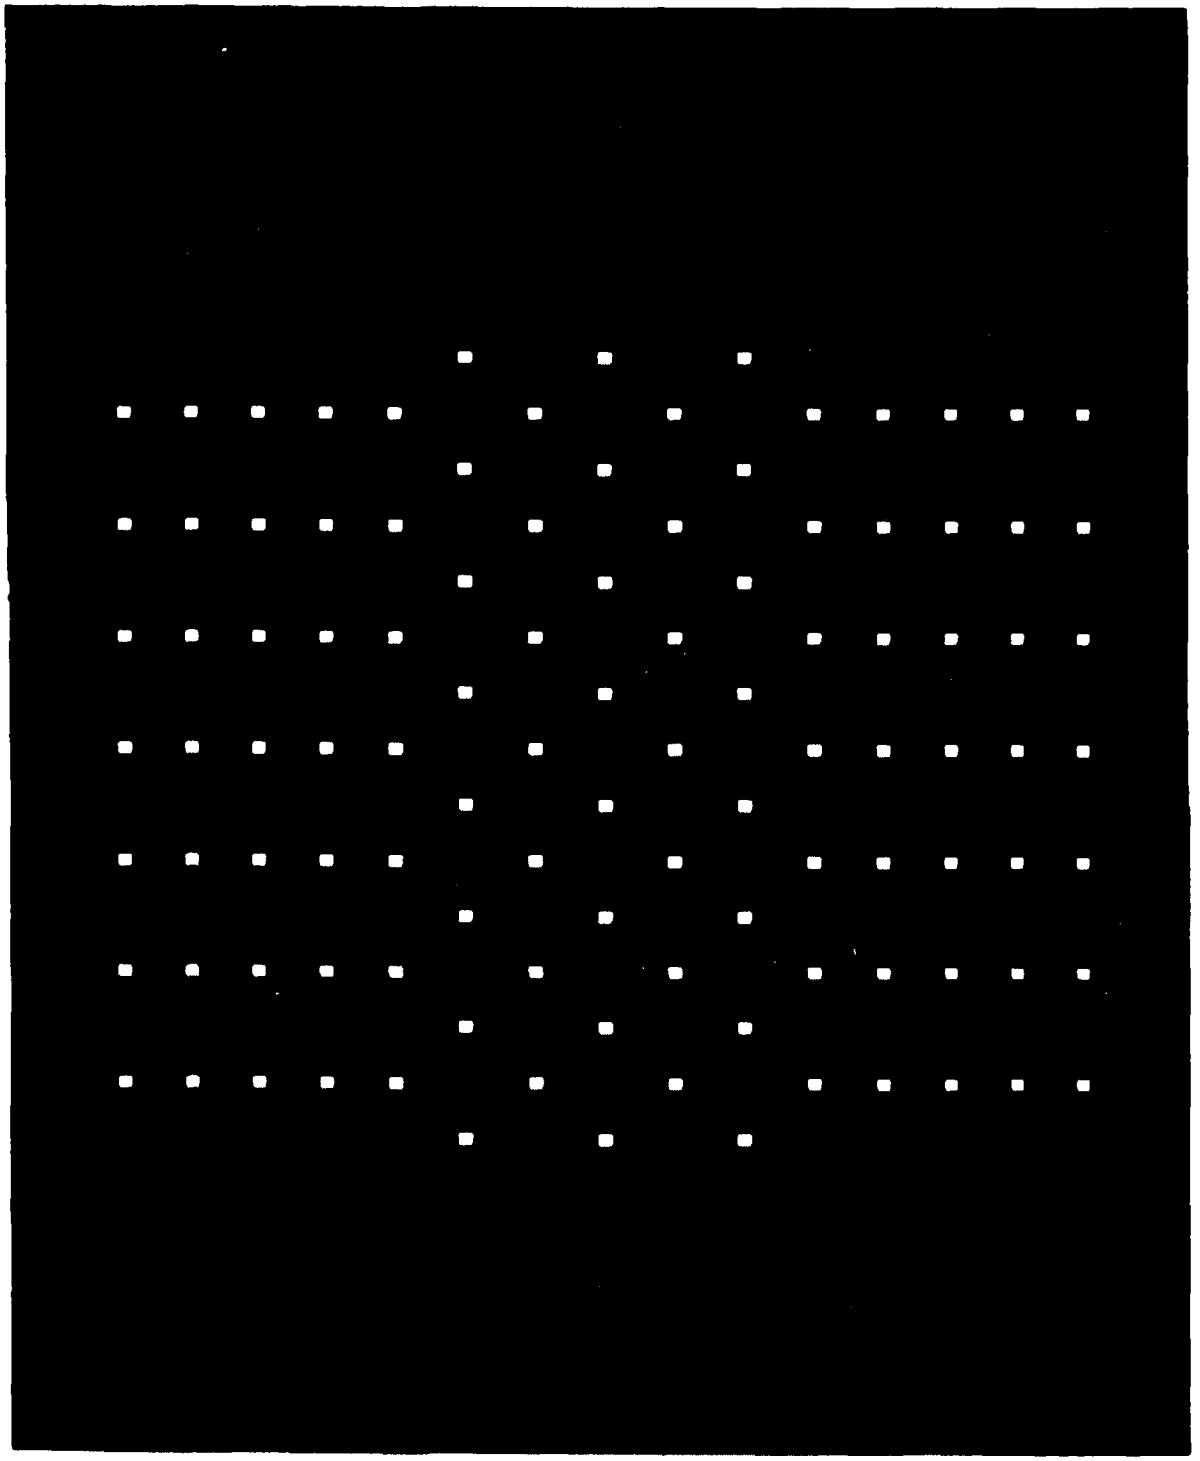

Figure 29

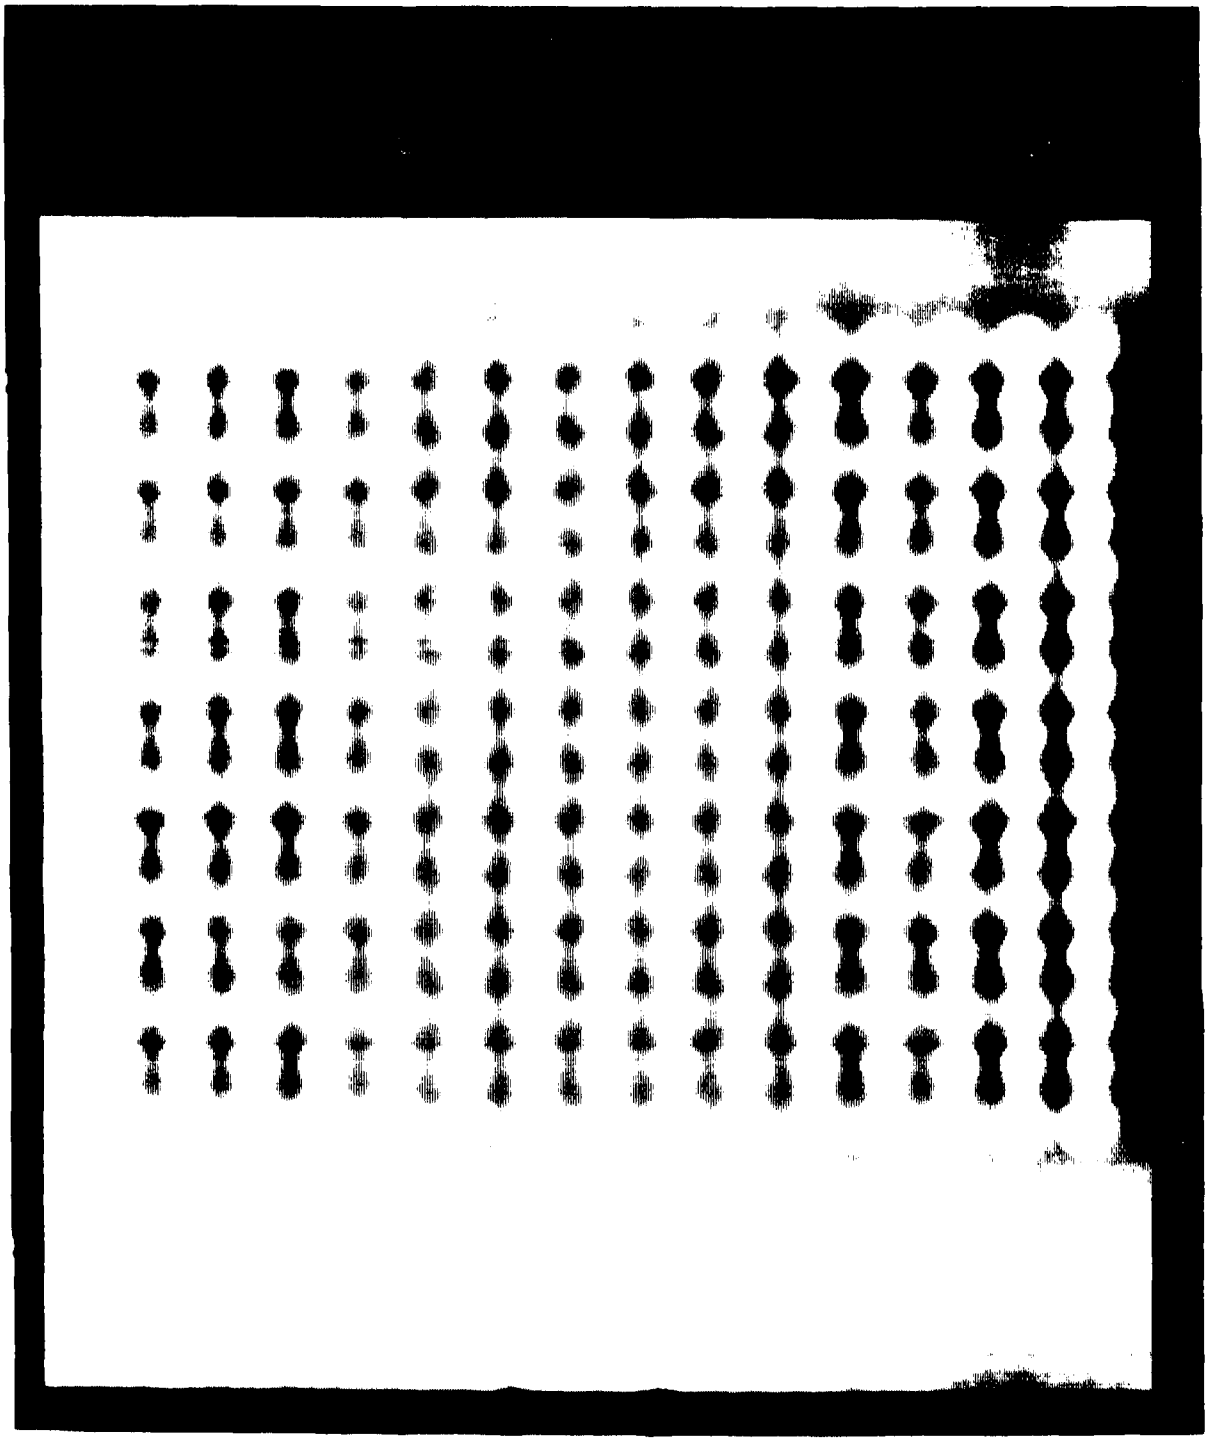

Figure 30

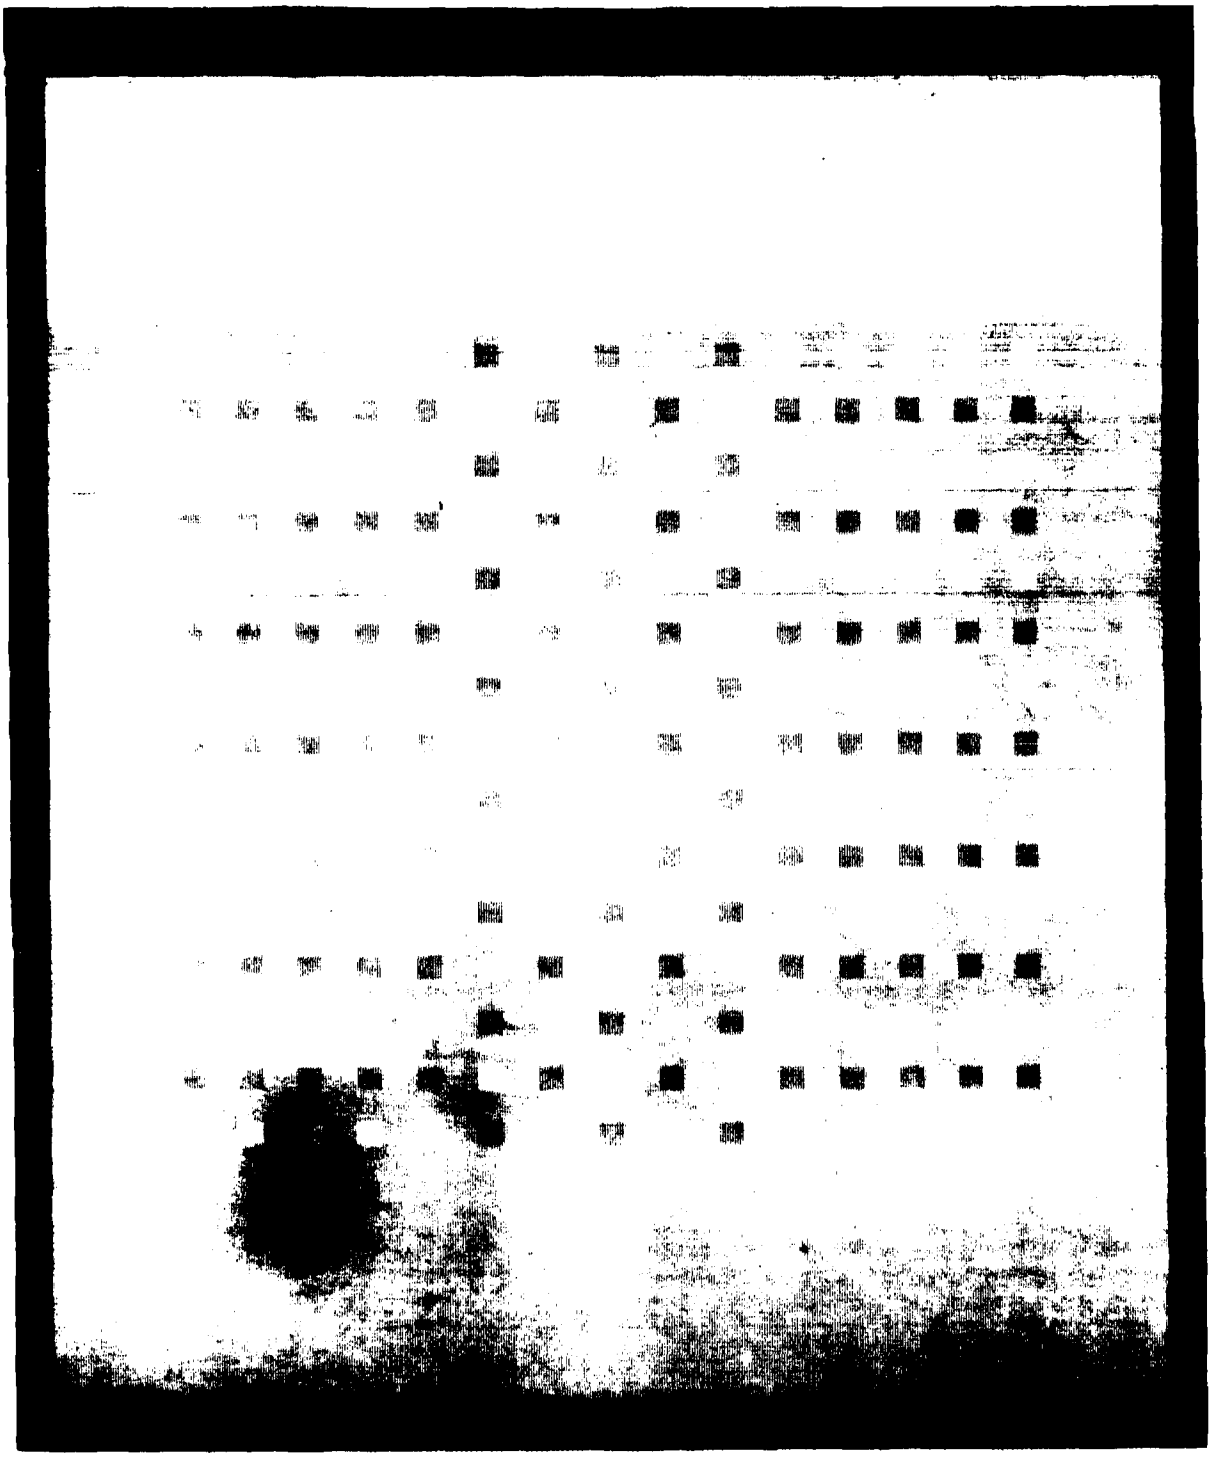

Figure 31

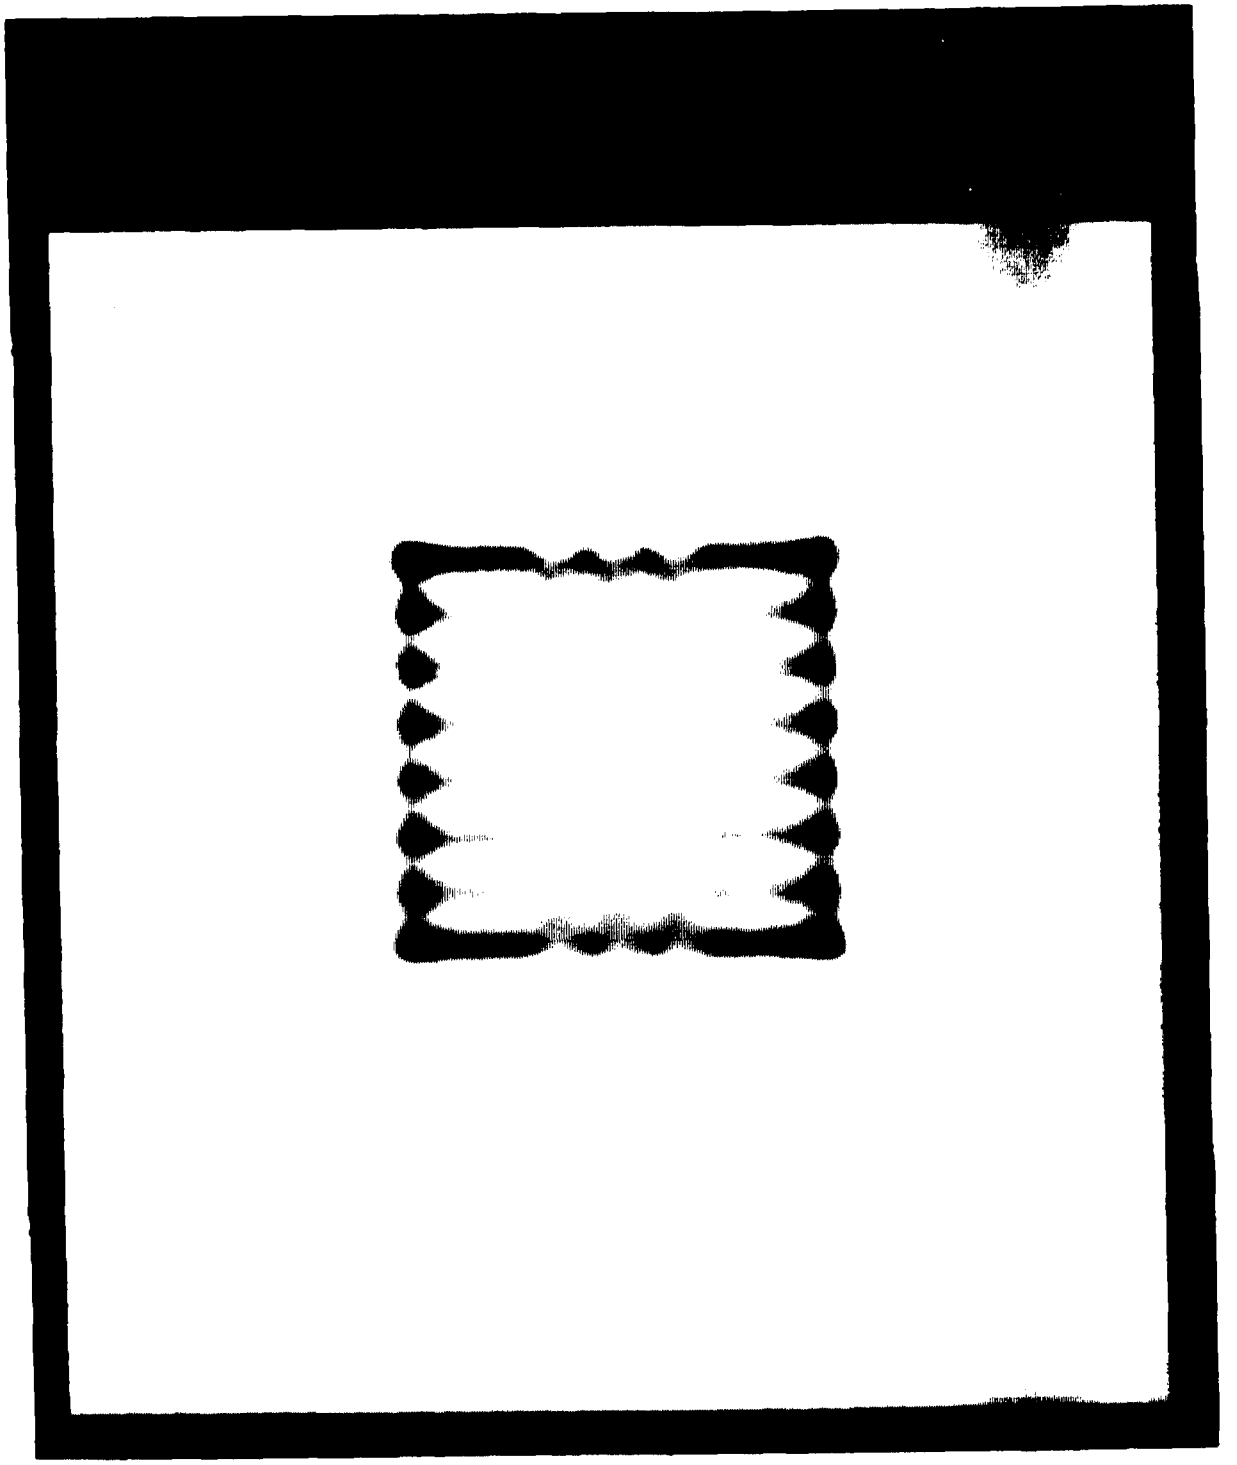

Figure 32

low frequency channels would signal vertically elongated blobs in the top and bottom regions.

Why does a difference in hue fail to produce tripartite segregation? We found that when the background and squares are made isoluminant, tripartite segregation occurs. As the background intensity is raised above the intensitites of the squares, tripartite segregation is lessened. Increasing the intensity of the background causes the achromatic receptive fields to respond. Because of the regularity of the tripartite pattern the receptive fields respond uniformly in all directions. Hue is a weak feature relative to energy. Hue differences will cause the squares in alternate columns in the top and bottom regions to link if not overriden by the achromatic receptive fields which respond equally in all directions.

2.5 Line Detection: Experiments 9-11

We have also studied how geometric properties affect preattentive grouping. Tripartite segregation can also occur in terms of emergent properties that result from the linking of texture elements (Beck et al. 1983). An example is the length of a line made up of broken segments. The paradigm we have used requires subjects to judge whether the line in the center of a stimulus was vertical or horizontal. Figure 33 shows an example of the displays used. The stimuli were presented on the Symbolics monitor and were flashed for 150 msec. They were viewed from a distance of 6 ft. with a pixel subtending 33 sec of arc. Ten subjects made 15 judgments of each stimulus. We recorded both reaction time and errors. The two measures closely agreed, and we shall report only the reaction times.

Experiment 9. Figures 34 and 35 show the variables investigated. One variable was element orientation. Figure 34 shows a square (6.6 min on a side) which has no orientation, (2) a weakly oriented rectangle (8.8 min x 5 min) and (3) a moderately oriented rectangle (10 min x 4.4 min). A second variable was the alignment of the element edges with the grouping axis. Figure 35 shows a blob stimuli derived from the bar stimuli by adding and subtracting pixels from the boundary. A third variable was the arrangement of the elements in a line. Figure 36 shows the arangement in which the elements (the weakly oriented rectangle) were arranged so that their axes were aligned (collinear) with the line. Figure 37 shows the arrangement in which the direction of the line is orthogonal to the orientation of the bars. Figure 38 shows the collinear arrangement of the weakly oreinted blob. The corresponding bar and blob stimuli were indistinguishable when they were convolved with a Gaussian with a sigma of 2 pixels. Differences in the grouping of the bars and blobs cannot be xplained in terms of low frequency energy differences.

Figure 39 show the mean reaction times. Both bar and blobs were detected more quickly when oriented [$F\ (1, 9) = 15.2\ p < .01$]. There was also a significiant interaction between contour alignment and arrangement [$F\ (1, 9) = 5.7\ p < .05$]. Bars were detected more quickly than blobs when the direction of the bars and blobs was the same as the direction of the line (collinear arrangement). On the other hand, blobs were detected more quickly when the direction of the bars and blobs were orthogonal to the orientation of the line (orthogonal

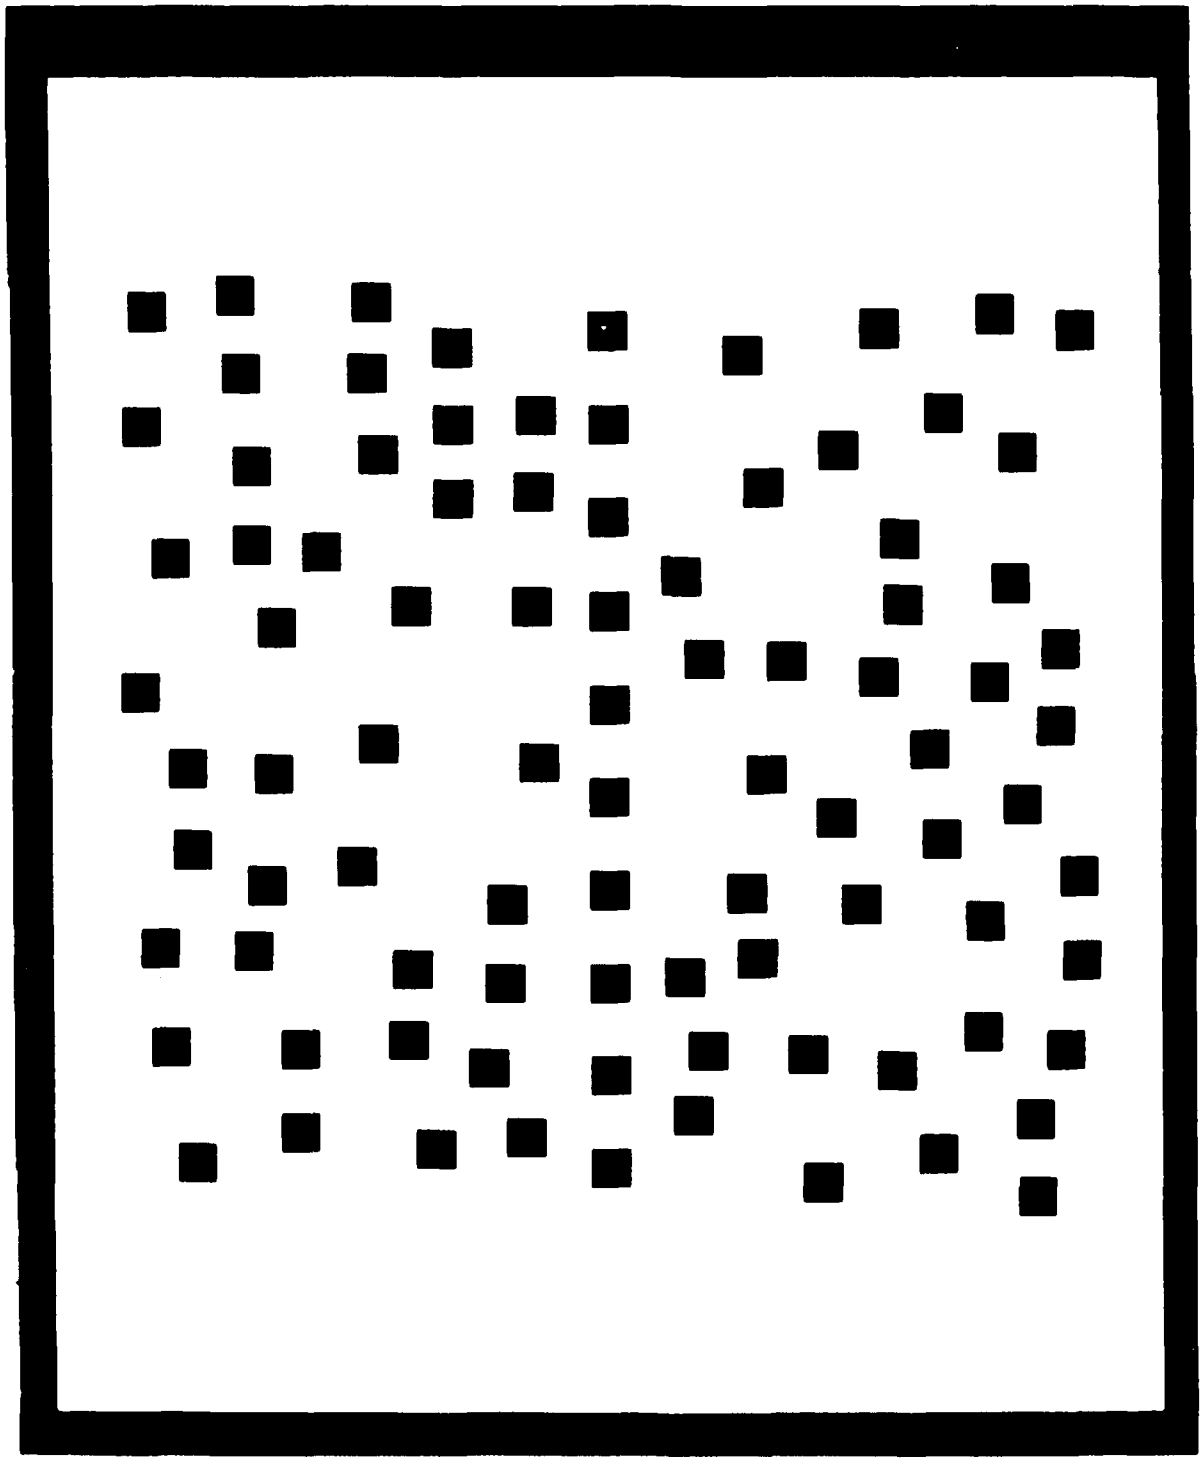

Figure 33

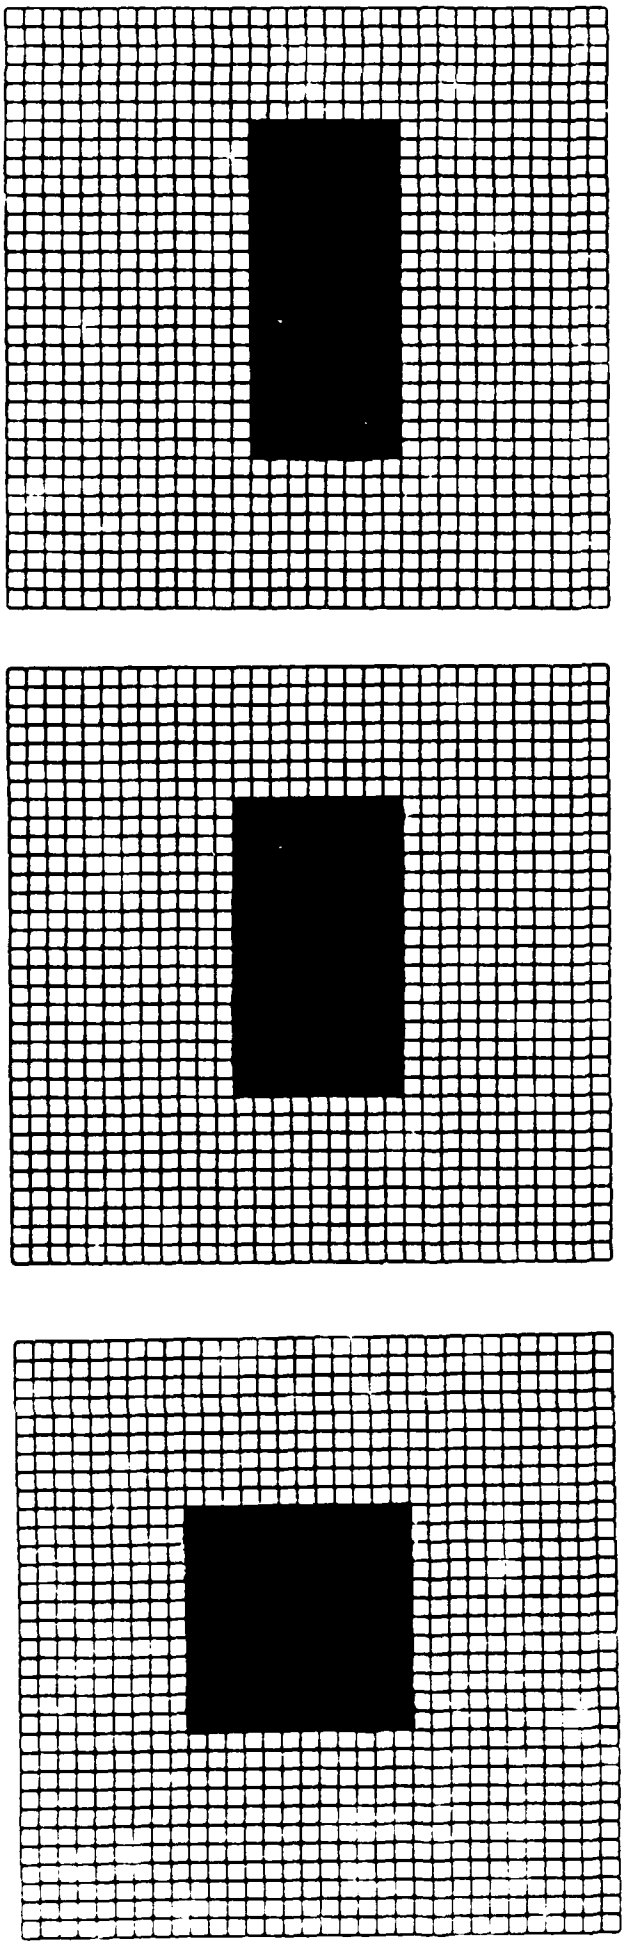

Figure 34

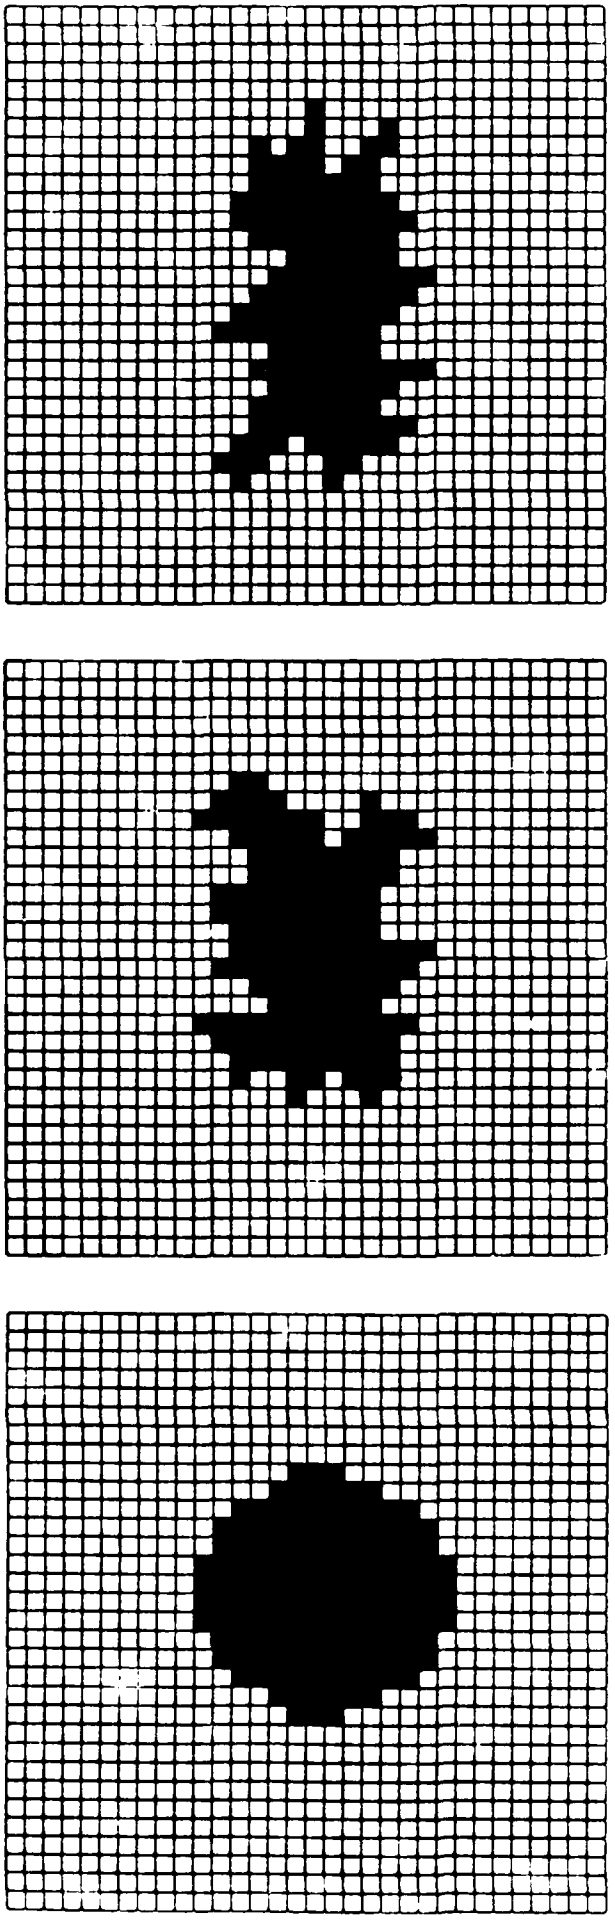

Figure 35

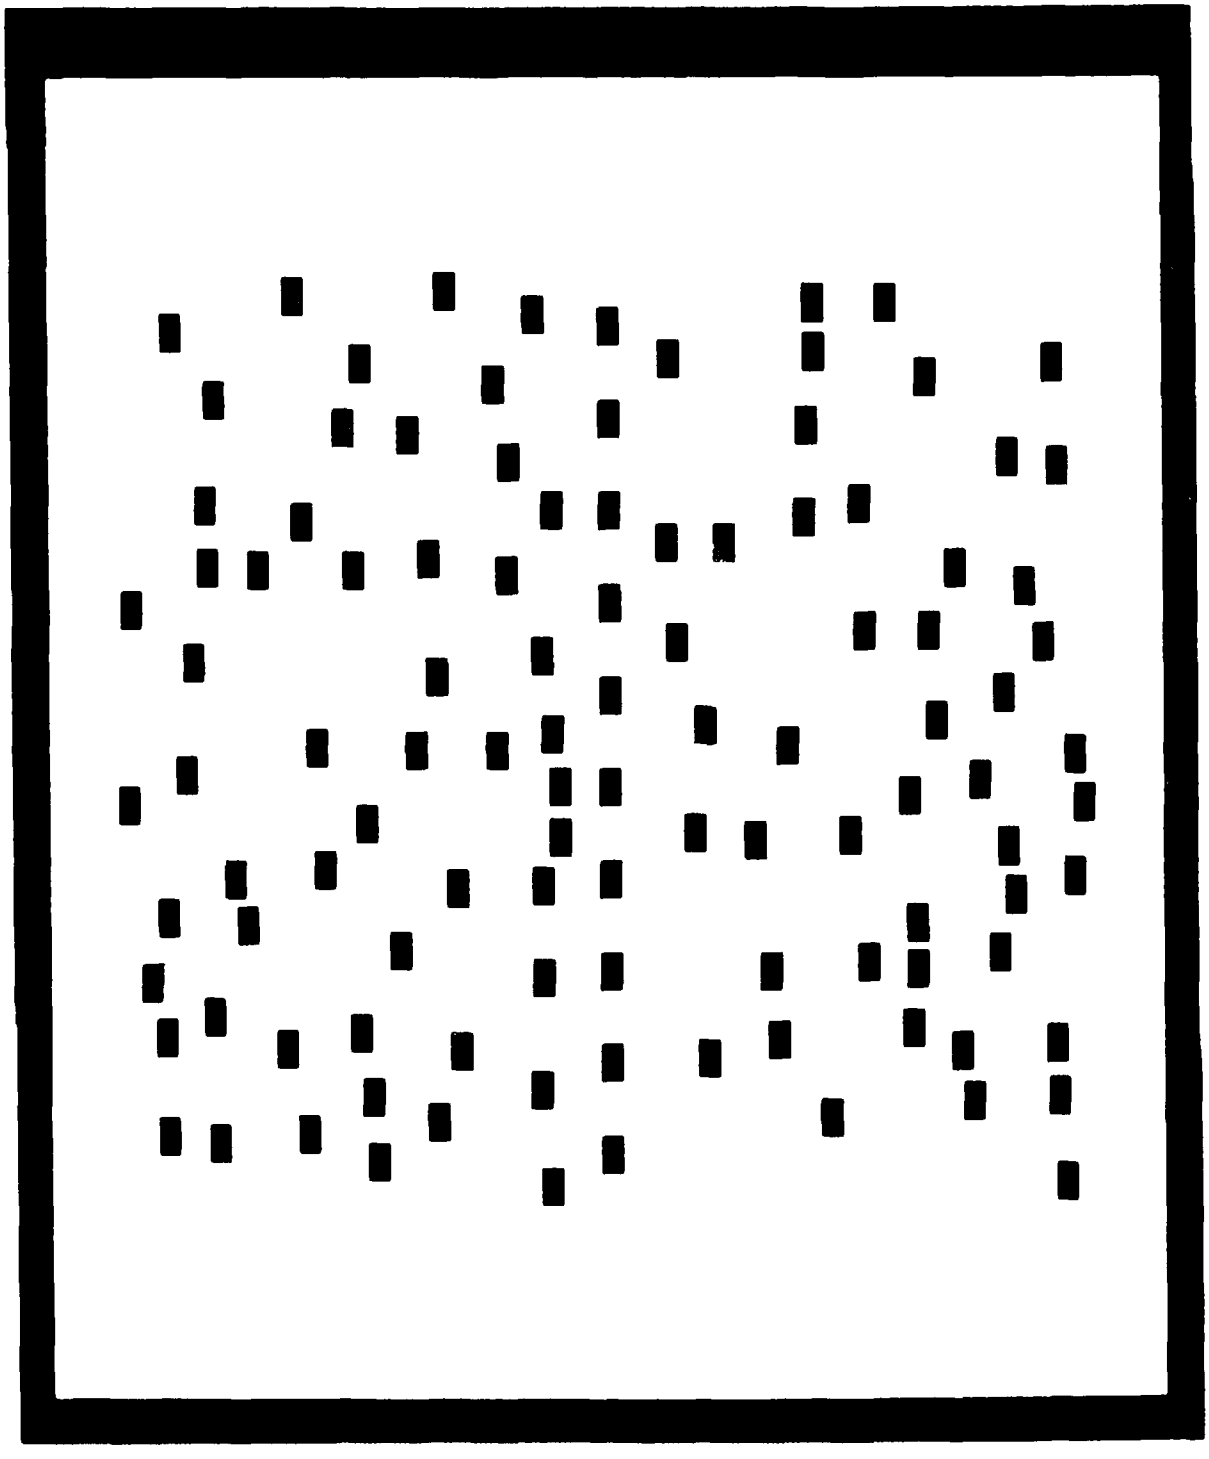

Figure 36

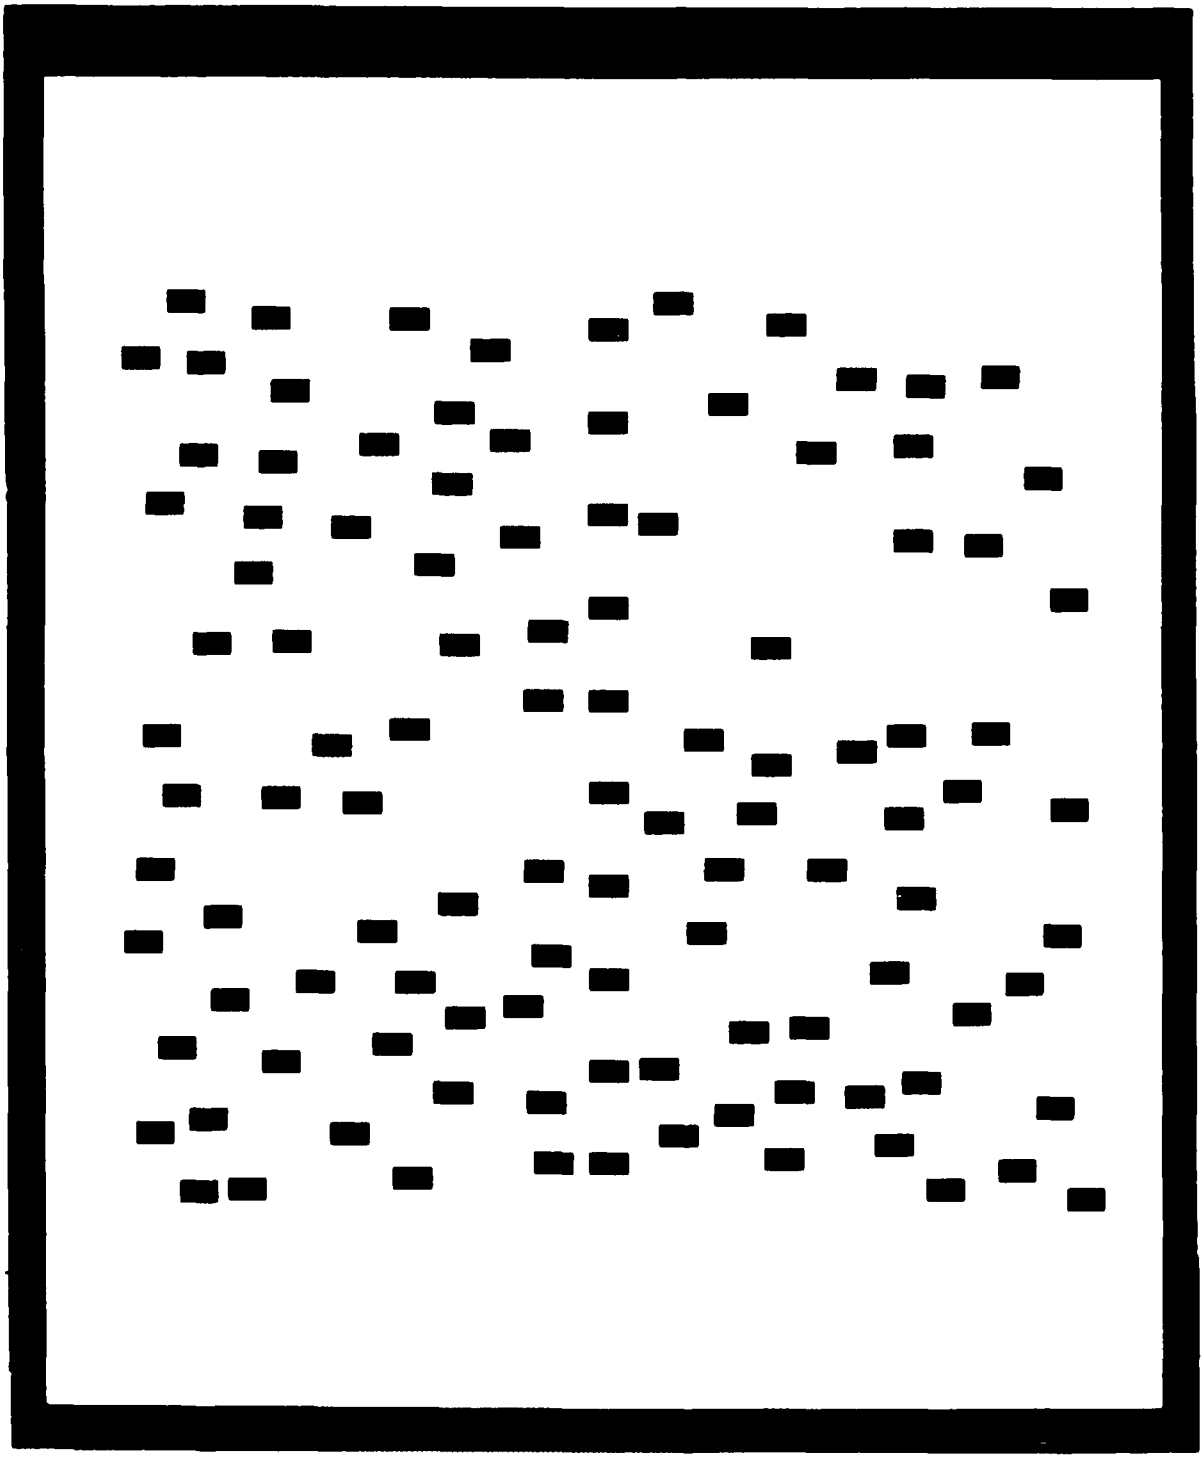

Figure 37

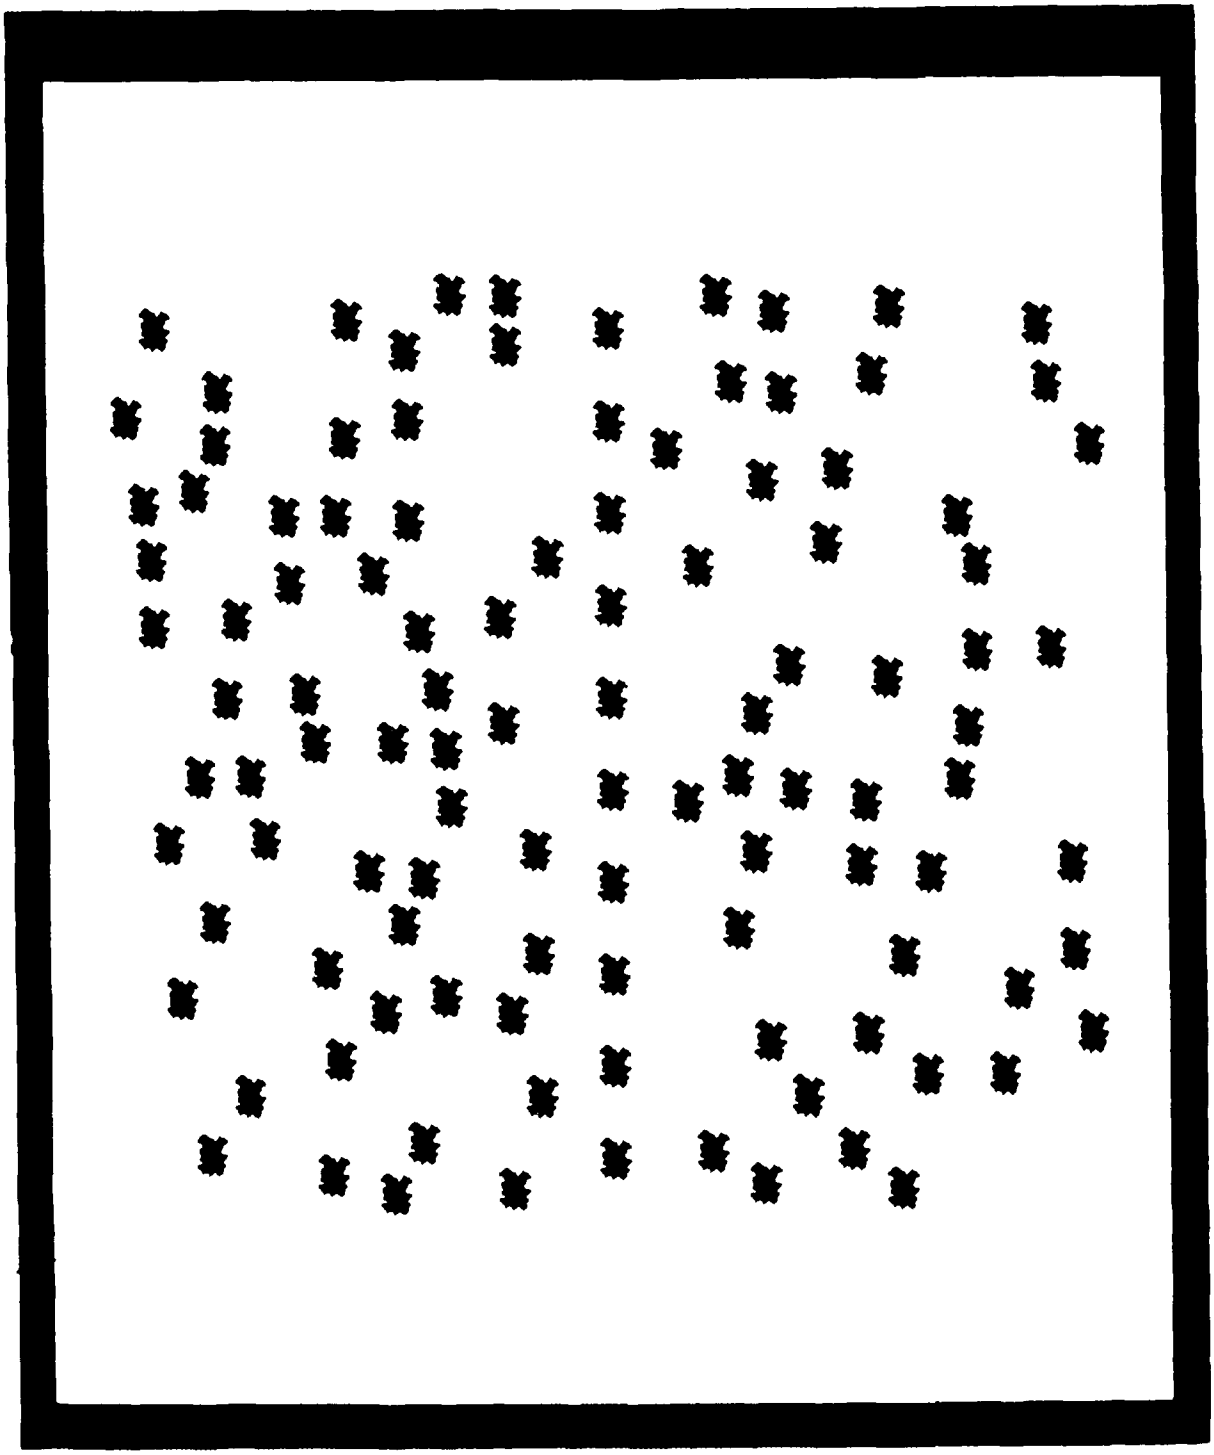

Figure 38

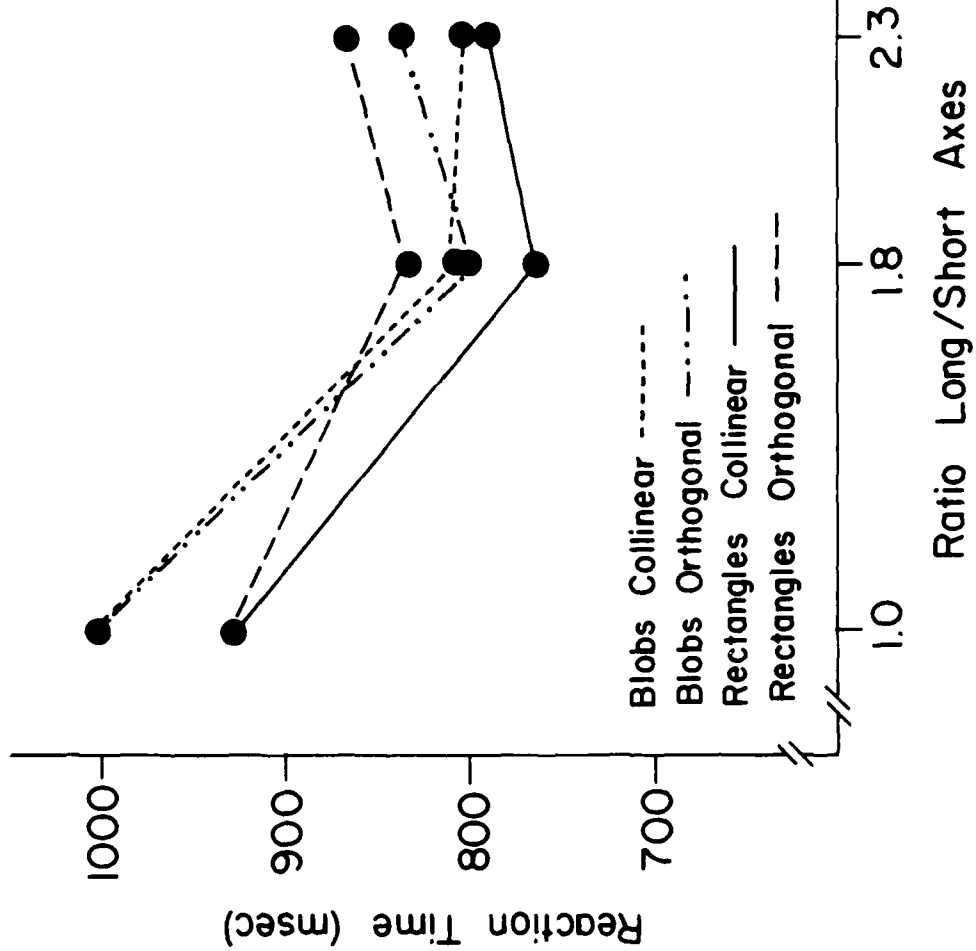
Reaction Time (msec)
1000
900
800
700
Blobs Collinear
Blobs Orthogonal
Rectangles Collinear
Rectangles Orthogonal
1.0
1.8
2.3
Ratio Long/Short Axes

Figure 39

arrangment). Orientation and smoothness of contour facilitated grouping the individual bars into a long bar when the orientation of the bars was collinear with the direction of the line.

Experiment 10. The variable of contour alignment was further investigated in Experiment 10. Figure 33 shows a stimulus in which the squares in the line are aligned. Figure 40 shows a stimulus in which the squares in the line have been laterally displaced by 25 percent of their height (four pixels). Experiment 10 compared the detection of a line when the line was composed of either squares or circles in which the circles and squares were aligned and in which they were laterally displaced by 12.5, 25, 37.5, and 50 percent of their height. Figure 41 shows the mean raction times. The reaction times increased with lateral displacement [$F(3, 27) = 29.4$ $p < .01$], and the reaction times for circles were significantly greater than for squares [$F(1, 9) = 10.5$ $p < .01$]. This confirms again the importance of contour in the grouping direction.

Experiment 11. Figure 42 show that mean reaction times were unaffected by stimulus scaling. That is, when the square size, lateral displacement, and the square separation were increased proportionally, the reaction time remained approximately constant. This indicates that the relevent variable is the relative orientation of the squares. The decrease in reaction time with stimulus scaling when the squares were collinear suggests that size is more important than separation for linking contours. The relative orientation of the squares also remains the same when the size of the squares are kept the same if the lateral displacement covaries with separation. However, under these conditions the reaction times increased. The grouping of squares is a function of their size and separation. Experiments 9 through 11 suggest that texture segregation can occur as a result of processes which link the contours of texture elements to produce new emergent properties.

2.6 Conclusions

The experiments present evidence for a two component theory of texture segregation. The first component is an operator sensitive to the outputs of spatial frequency sensitive mechanisms. This operator segregates regions according to differences in contrast sign and differences in low spatial frequencies. There is no interaction between positive and negative contrasts, and strong segregation can occur even when there is only a small difference between the high and low intensity squares when the intensity of the background is in between. If the contrast sign is the same, texture segregation occurs only if the shapes of the distributions of the low spatial frequency mechanisms responding differ. That is, high frequency differences are not sufficient to cause segregation in the presence of contradictory information from low frequency mechanisms. Thus, size and contrast can cancel. A small size and a high contrast that stimulate the same low frequency mechanisms as a large size and a weak contrast fail to cause texture segregation. Equal size squares that differ in contrast magnitudew will not cause texture segregation unless the contrast difference is very large. Differences in contrast magnitude unless very large do not change the shape of the distribution of low spatial frequency mechanisms responding. Thus,

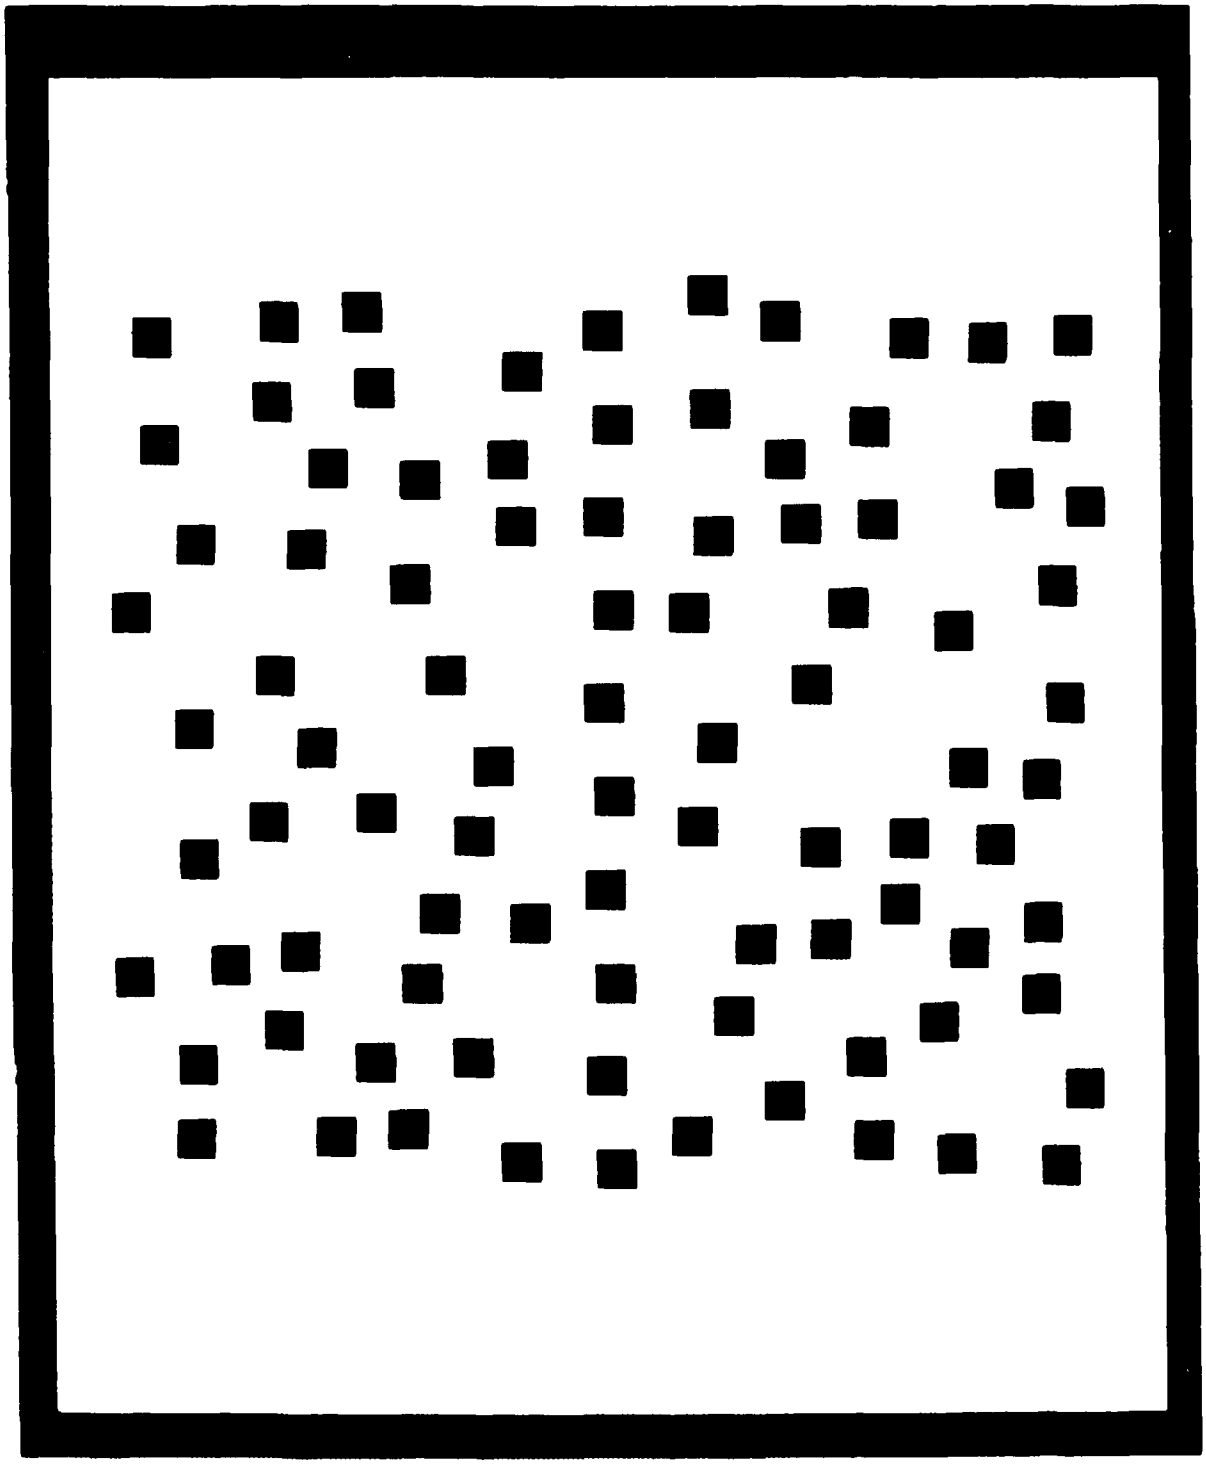

Figure 40

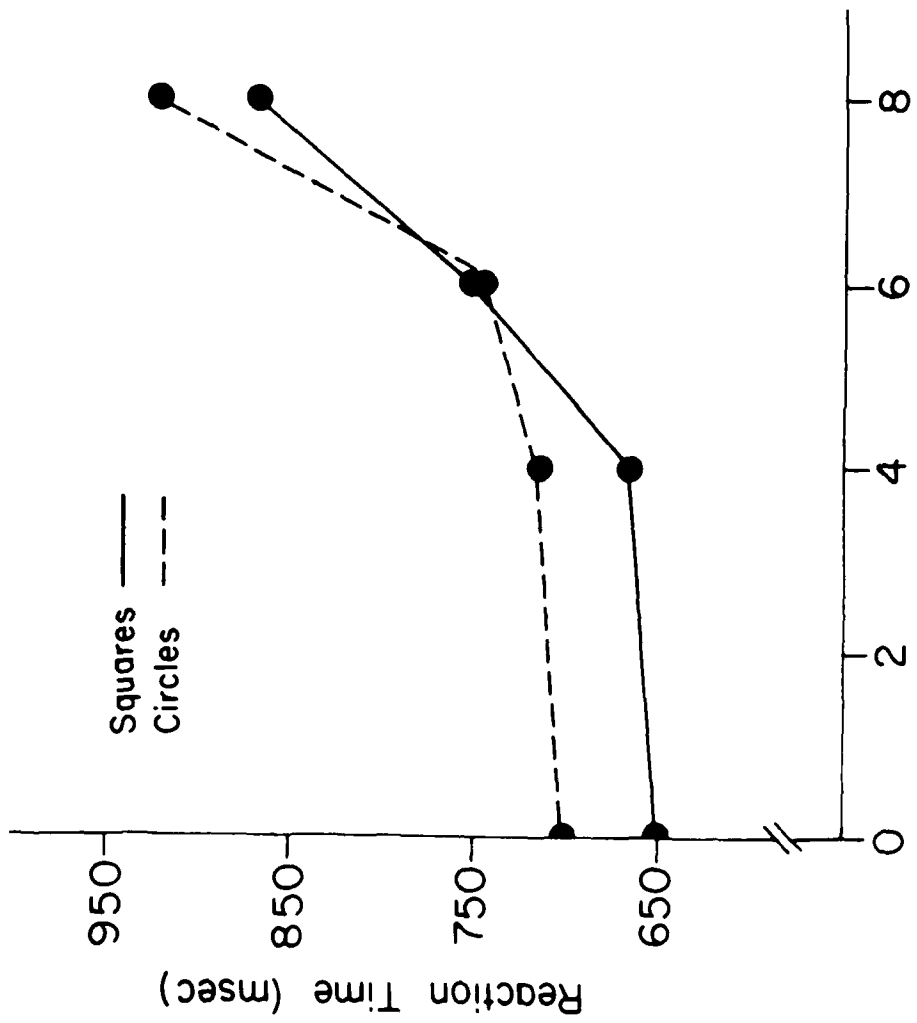
Squares
Circles
950
850
750
650
Reaction Time (msec)
0
2
4
6
8
Lateral Displacement in Pixels
(1 pixel = 33 sec of arc)

Figure 41

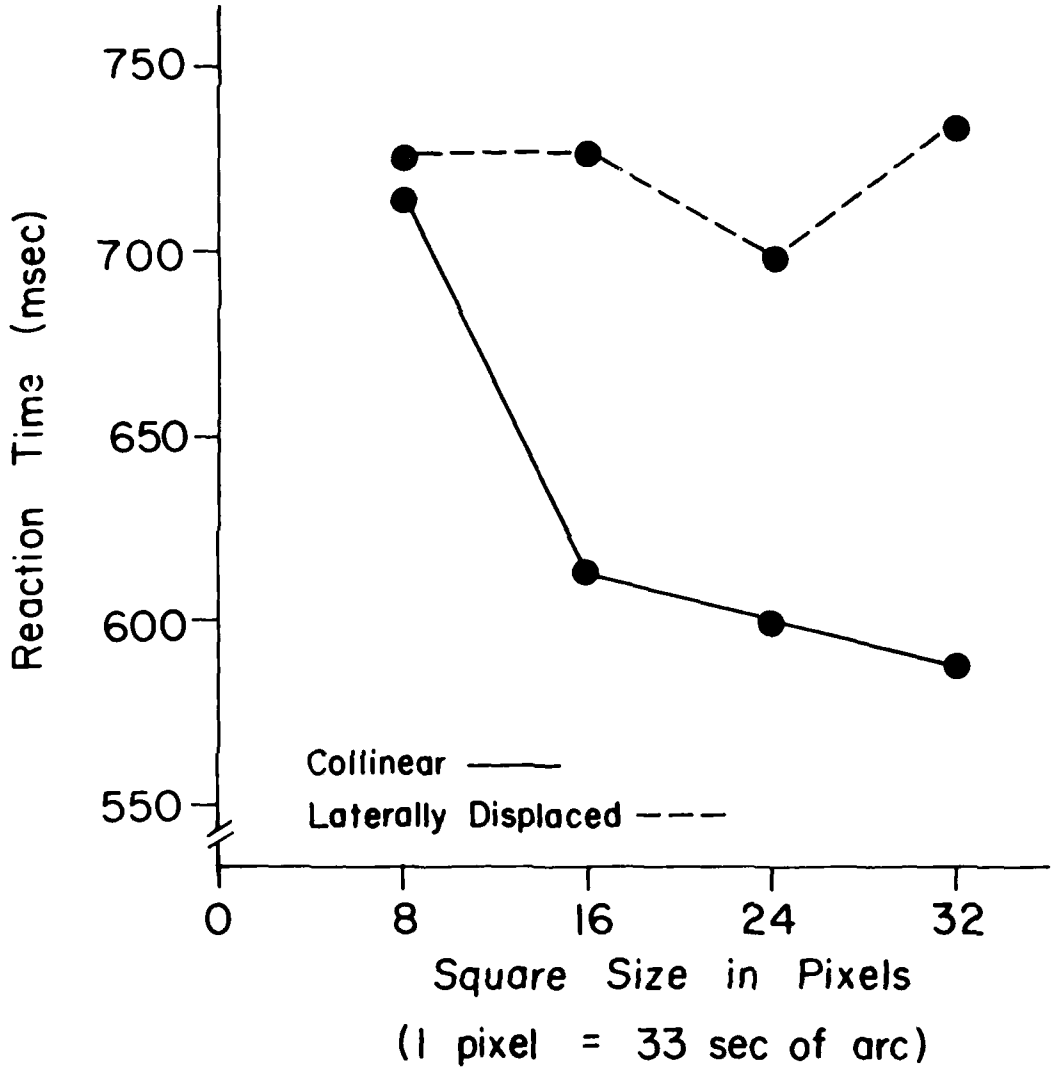
750
700
650
600
550
Reaction Time (msec)
Collinear
Laterally Displaced
0
8
16
24
32
Square Size in Pixels
(1 pixel = 33 sec of arc)

Figure 42

readly discriminable differences in lightness fail to cause texture segregation. Texture segregation and the discrimination of lightness involve different mechanisms. Differences in contast magnitude do cause texture segregation with very close spacing because a new element, bar detectors, are now stimulated. Hue is a weak feature relative to contrast. That is, hue differences are not sufficient to cause texture segregation in the presence of contradictory information from contrast mechanisms. The second component is an operator that aggregates texture elements on the basis of geometric descriptors such a aspect ratio, alignment, contour smoothness etc., and then segregates the regions on the basis of "emergent" properties. The effects of these variables can not be explained in terms of differences in the response of low spatial frequency mechanisms. Th linear organization appears to be derived from a nonlinear operator that connects nearby localized spatial tokens.

Our present work is concerned with defining more precisely the shape differences in the outputs of the low frequency mechanisms that causes texture segregation, the necessary and sufficient conditions for texture aggregation to occur, and the interaction (if any) between the spatial frequency and aggregation mechanisms in texture segregation.

3. DETECTING STRUCTURE BY SYMBOLIC CONSTRUCTIONS ON TOKENS

3.1 Introduction

In this section we report evidence that geometric organization in texture emerges by constructions on symbolic tokens that represent the constituent elements [Stevens & Brookes 1986]. In a dot pattern, for example, the individual dots are represented as discrete items and given attributes including color, size, and contrast. The organization in the dot pattern emerges by operations on these items and not the original retinal intensity distribution.

The notion of discrete grouping items or tokens was tacit, by and large, in the Gestalt demonstrations of similarity grouping [Wertheimer 1923; Koehler 1929; Koffka 1935]. Later, it was specifically proposed that groupings involved "place markers" [Attneave 1974] or "place tokens" Marr [1976, 1982], which individually carry information about position and attributes such as contrast, color, size and orientation (see also Ullman's [1979] grouping tokens for motion correspondence). A primitive operation on place tokens would be to group tokens pairwise, with the pairing represented by a "virtual line" [Attneave 1974; Marr 1976, 1982; Stevens 1978]. Similarly, Caelli and Julesz [1978] discuss "local dipoles" between neighboring dots in texture discrimination. While virtual lines do not manifest apparent contrast (as do subjective contours) they nonetheless seem to make orientation and separation visually explicit. Attneave [1955, 1974] has found evidence for position and orientation judgements mediated by place tokens, and recently Beck and Halloran [1985] have suggested that virtual lines might play a role in vernier acuity judgments. It is not clear from such experiments, however, whether the position markers that seem to mediate attentive, foveal judgments of relative position and orientation under scrutiny are the same "place tokens" that have been proposed for early visual processing of texture. There are alternative models for how groupings might emerge in texture that do not require explicit position markers.

A common proposal is that elongated receptive field mechanisms such as simple cells detect collinear organization in discrete patterns, without need for explicitly marking the constituent items. Instead, the receptive field would compute some spatial average by means of an arrangement of excitatory and inhibitory subfields, and the structure would be revealed by means of the receptive field's orientation selectivity. Hence while groupings seem to naturally require some explicit representation of the constituent elements, receptive field mechanisms seem able to detect their alignment without distinguishing the individual elements.

Arguing the case for place tokens would benefit from a more precise definition of what might constitute a place token. Thus far, however, the argument for place tokens has been largely intuitive, with little discussion regarding what specific intensity events or features might define place tokens. Hence the approach we have taken is to capture the central notion with the following dichotomy: either the perceptual grouping of discrete items either distinguishes them as individual entities, or not. The former implies place tokens, the latter reduces to, in effect, spatially blurring the discrete pattern into a continuous distribution. Specific proposals

involve summation of dot energy within simple-cell receptive fields (see below). Evidence against such schemes would constitute indirect evidence for place tokens. We will show the inadequacy of spatial blurring schemes in general, and sketch an argument supporting this finding, and specifically show cases where the simple cell model is inadequate. That is, various patterns can be constructed in which the linear organization is visually apparent but for which elongated receptive fields of the simple cell variety would be ineffective. Our interpretation is that the visual system in those cases extracts the organization by means of place tokens. It should be stressed that the evidence does not rule out the simple cell model for detecting dot pairings and collinearity, but rather shows where it is an inadequate explanation. Our goal has been to examine the primitives of a class of perceptual grouping, and in so doing to pry apart various ideas about neural mechanism from the more basic issues of what computations must be performed.

3.2 The Place-Token Hypothesis

At an early stage of processing, the visual system generates an array of local intensity descriptors, each of which describes the intensity change (e.g. bar or edge of given contrast, orientation, and so forth) occurring at the corresponding retinal position. Marr [1982] refers to this as the raw primal sketch (RPS). It corresponds to the stage at which LGN X-cell input (primarily) is interpreted both spatially and across scales in order to localize and describe intensity features. The RPS is local and unarticulated — the constituent assertions (e.g. of edge or bar) are not organized into larger ensembles, and their spatial arrangement is unknown (e.g. the local connectivity among adjacent edge or bar segments remains implicit). Using Marr's [1982] terminology, the full primal sketch (FPS) refers to the image description at which structure is explicitly represented and the local intensity changes are organized in a manner that reflects their physical causes and arrangement. The FPS contents are organized by processes of both aggregation and segregation — aggregated so that elements having a common physical origin are associated, and segregated so that physically distinct regions are distinguished.

The task of imposing organization on the RPS is prodigious, and seemingly open-ended (in that even simple spatial relations such as global connectedness, closure, and inside-outside require considerable computation for even small collections of elements, and would be combinatorially prohibitive to compute across the RPS [Ullman, 1984]). It is therefore not clear which basic organizing processes are performed routinely to generate the FPS from the RPS, but it is probable that certain, computationally tractable spatial relations, particularly collinearity and parallelism, are detected in the RPS over at least moderately global spatial extents. Collinearity and parallelism are singled out here as representative of simple geometric relations that are extracted early in vision, and which constitute part of the FPS. They are likely computed preattentively since they are compellingly apparent in patterns that are presented briefly under experimental conditions such that the given stimulus structure cannot be predicted.

3.2.1 Constructs for Indexing and Representing Selection

The familiar grouping phenomena associated with dot patterns likely reflect organizing processes that impose organization on the RPS. When a row of dots emerges as a dotted line, for example, the visual system has grouped those dots into a new "emergent" organization. The question we have posed thus far is whether the tangents to the dotted line along its length are "detected" by receptive fields or "constructed" as local pairings between adjacent tokens. More generally, these perceptual grouping processes must impose organization that accurately reflects the structure of the visual scene. A simple case concerns adjacent edge segments that are collinear, contiguous, and similiar in contrast, color and other properties very likely correspond to adjacent physical events. Their common association is clear, and various means are readily suggested for tracing the continuous curve. But the correct grouping is less obvious when edge segments are not contiguous, as when interrupted by interposition of other opaque surfaces, or when the contrast vanishes along a curve due to variations in the background intensity. The general expectation is that the visual system determines the correct association partly on the basis of similiarity of various attributes. The argument, of course, is that physically related intensity changes are likely to be similiar along various dimensions (both spatial, such as orientation, scale, sharpness, direction of motion, and contrast distribution, and non-spatial, such as color and intensity). This sketch serves our immediate purposes; we can now relate the place token hypothesis to the problems of imposing organization on the RPS.

To impose any structure on a collection of elements on the basis of their geometry and attribute similiarity, requires two basic computational abilities:

i) A means to address or index elements by position and by attribute,

ii) A means to represent the selected subset of elements.

In other words, an access mechanism is need to index individual elements out of a set. The particular access mechanism can remain unspecific for now; important is the notion of filtering or selecting a subset of a collection according to some criterion or criteria. For instance, it seems reasonable to expect the visual system is able to access edge segments of a specific orientation range (say, roughly vertical) within a specific spatial region, or to select any and all elements moving to the right, and so forth. (We will discuss selection more below.) In addition to an access mechanism that allows extraction of a given subpopulation, it is also necessary to represent these distinguished elements as a cohesive entity. In concrete terms, when a line of collinear dots appears to stand out from the random arrangement of background dots, there must have been a selection step, to extract collinear dots, and a new construct introduced to represent them as a single entity. The spontaneous organization of a dot grid into parallel rows or columns cannot be accounted for merely in terms of receptive field mechanisms (even assuming such mechanisms do detect the individual rows or columns). The local evidence for each individual row or column must be extracted and grouped into distinct ensembles.

There is need, therefore, for a means for accessing elements by attribute and position and for addressing neighboring places where similar intensity changes occur. We suggest that (i) place tokens provide this facility, and that (ii) virtual lines make explicit the spatial relationship between similar tokens.

3.2.2 The Scales of Intensity Change and of Structure are Independent

A new perspective on structure detection is provided, we believe, by the following premise: that the scale of geometric structure is independent of the scale of the intensity changes that comprise it. Thus, for example, global organization is not necessarily carried by the more global features. Line segments detectable at the finest scales of resolution (present in only the very high spatial frequencies) might be arranged collinearly or in parallel striations that is detectable only by examining these elements across a spatial scale an order of magnitude larger than their component spatial frequencies. We see two distinct tasks, therefore, in detecting intensity changes across spatial scales, and in detecting structure across spatial scales. It is clear that intensity changes may occur at several scales within a given spatial area — minute edges and markings might be superimposed over large-scale intensity features. Likewise both fine-scale and larger-scale structures might superimpose within an image, such as found in the texture of a herringbone fabric. At a small scale are thin lines corresponding to the individual fibers, at a slightly larger scale the parallel diagonal striations characteristic of the herringbone is present, and at a larger scale their vertical organization into columns is apparent. At a still larger scale one might observe folds and creases across the fabric.

As we have found, the extraction of a structure at any scale seems largely independent of the scale of the individual elements. We will use that observation as part of our case for the place token hypothesis, as it is difficult to reconcile such results with the simple cell model. But more generally, there are geometric relationships that require both acute sensitivity to position information and scale independence (which usually act in opposition). The resulting apparent organization is not captured by any single scale of image description. Rather it appears necessary to generate distinct structural assertions that make explicit or summarize the local geometry. We thus return to the idea that local geometric organization emerges by synthesis, not by detection. The general conclusion we draw is that structure computations need to be considered less from the point of view of the familiar "feature detectors" such as simple cells, such as correlations or cooperative computations that sharpen their effective orientation tuning curves. We will close this discussion with some suggestions.

3.3 Alternative Models

If the constituent elements in some geometric grouping are not explicitly marked, their local arrangement must be detected from their averaged spatial distribution, usually phrased in terms of blurring (low-pass filtering) or energy summation within elongated receptive fields. For example, a sufficiently-closely spaced pair of dots, or a chain of collinear dots, has a power spectrum similar to that of an isolated line segment for spatial frequencies less than $1/s$, where

s is the dot spacing. The dotted line is thus roughly equivalent to a continuous line having equal total energy as stimulus to a linear summation mechanism, such as the receptive field organization of an even-symmetric, "bar-detector", simple cell.

Glass proposes that the local orientation is derived by correlating the activity of simple cells over small neighborhoods [Glass 1969; Glass & Perez 1973; Glass & Switkes 1976; Glass 1979]. Zucker [1983] similarly proposes a cooperative computation whereby the broad orientation tuning curves of individual receptive fields can be sharpened by combining the outputs of individual cells over local neighborhoods. By such proposals, simple cells whose elongated, bar-shaped receptive fields align with the dot pairs would respond more vigorously, on average, than those cells at other orientations, hence local correlation (or similar computations) of their activity would reveal the orientation of the dot pairs in each vicinity. Similarly, Caelli and Julesz [1978] suggest that linear arrangements of dots in texture are detected by neural units with elongated receptive fields applied to the retinal image, either "a single neural feature extractor of the Hubel and Wiesel type", or a unit that "measures the quasicollinearity of adjacent dipoles by combining single neural units of a retinal neighborhood with slightly different orientation sensitivity" [Caelli & Julesz 1978, p. 172; see also Caelli *et al.* 1978; Julesz 1981]. Recently Prazdny [1984] also suggested that the dot pairings in Glass patterns are detected from "...measurements in the spatial and energy domain rather than logical operations on symbolic descriptions".

3.4 Dot Pattern Phenomenology and Pitfalls

A deceptively simple dot pattern that has become popular for examining grouping phenomena is the Glass pattern [Glass 1969]. Glass patterns (figure 1) are constructed by superimposing onto a random dot pattern a copy that has been transformed, *e.g.* by scaling or rotation. Each dot and its transformed counterpart in the superimposed copy defines a dot pair. If the copy is scaled, say, a globally radial pattern emerges, where the dot pairs are all radially aligned. The dot pairs have been characterized as tangents or vectors, and it can be shown that the apparent overall organization emerges from the local orientation of individual dot pairs [Stevens 1978]. (A pure rotation or scaling would result in dot pairs of variable separation across the pattern, and as a result the apparent grouping would depend on where one fixated. This is easily avoided by generating Glass patterns with homogeneous displacements between corresponding dots in order to have the pairs equally apparent across the pattern [Stevens 1978].)

While the visual effect in the Glass pattern seems to arise from the detection of local dot pairings, there are clearly other factors than the dot pairs themselves contributing to the apparent organization, principally inhomogeneities in dot density that arise as artifacts of the process of generating a Glass pattern, as discussed below.

3.4.1 Large Scale Clusters and Density Inhomogeneities Dominate

When Glass patterns are used to test theories of dot grouping, there is a tendency to concen-

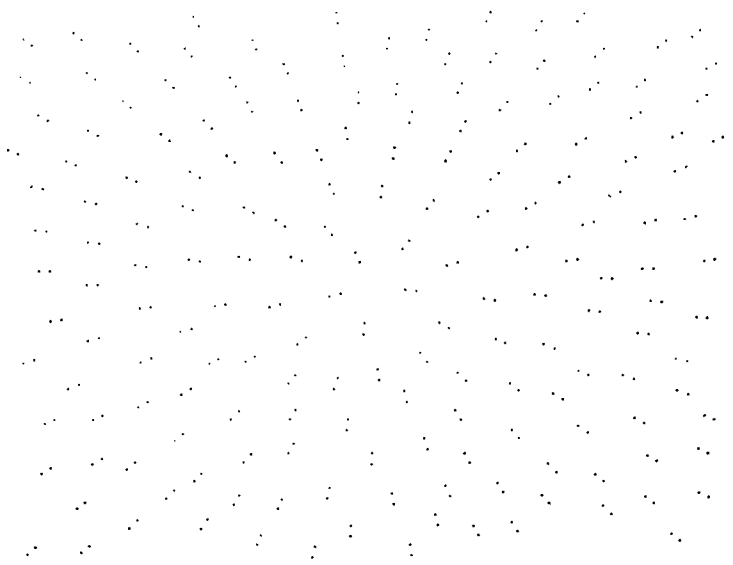

Figure 1. A radial Glass pattern based on a homogeneous density dot pattern, with homogeneous dot displacements.

---

trate on the subjective pairing of dots, and to attribute perceived global organization to the locally detected pairings or dipoles. Unfortunately, Glass patterns, by the way they are generated, tend to accentuate density inhomogeneities along the local transformation direction [Stevens 1978]. If n dots in the original dot pattern happen to align along the direction the copy will be translated, the resulting pattern will have $2n$ collinear dots. This results in chains of four, six, or more dots as well as the simple pairings of dots, if not controlled for. Furthermore, inhomogeneities in dot density, which appear as clusters or voids, are selectively enhanced in the direction of the local transformation by the same process, resulting in an eccentuation of the boundaries of the clusters and voids along the same path as the individual dot pairs. Figure 2a, for example, shows a conventional Glass pattern of moderate density and dot displacements, built from a random dot pattern. While in figure 2a the concentric organization seems carried by the dot pairings, in figure 2b the displacement is so large that the pairings are no longer apparent but nonetheless the overall effect is still apparent. Clearly the low spatial frequency components in this pattern are responsible for the apparent organization, and the dot grouping phenomena are secondary.

When dealing with Glass patterns, therefore, it is necessary to make dot density as homogeneous as possible to avoid clusters and voids. We use "homogeneous density", (constant nearest-neighbor distance) dot patterns to minimize these effects. The corresponding Glass pattern (figure 2c) presents clear dot pairings, which is largely extinguished when the corresponding dots no longer appear paired (figure 2d).

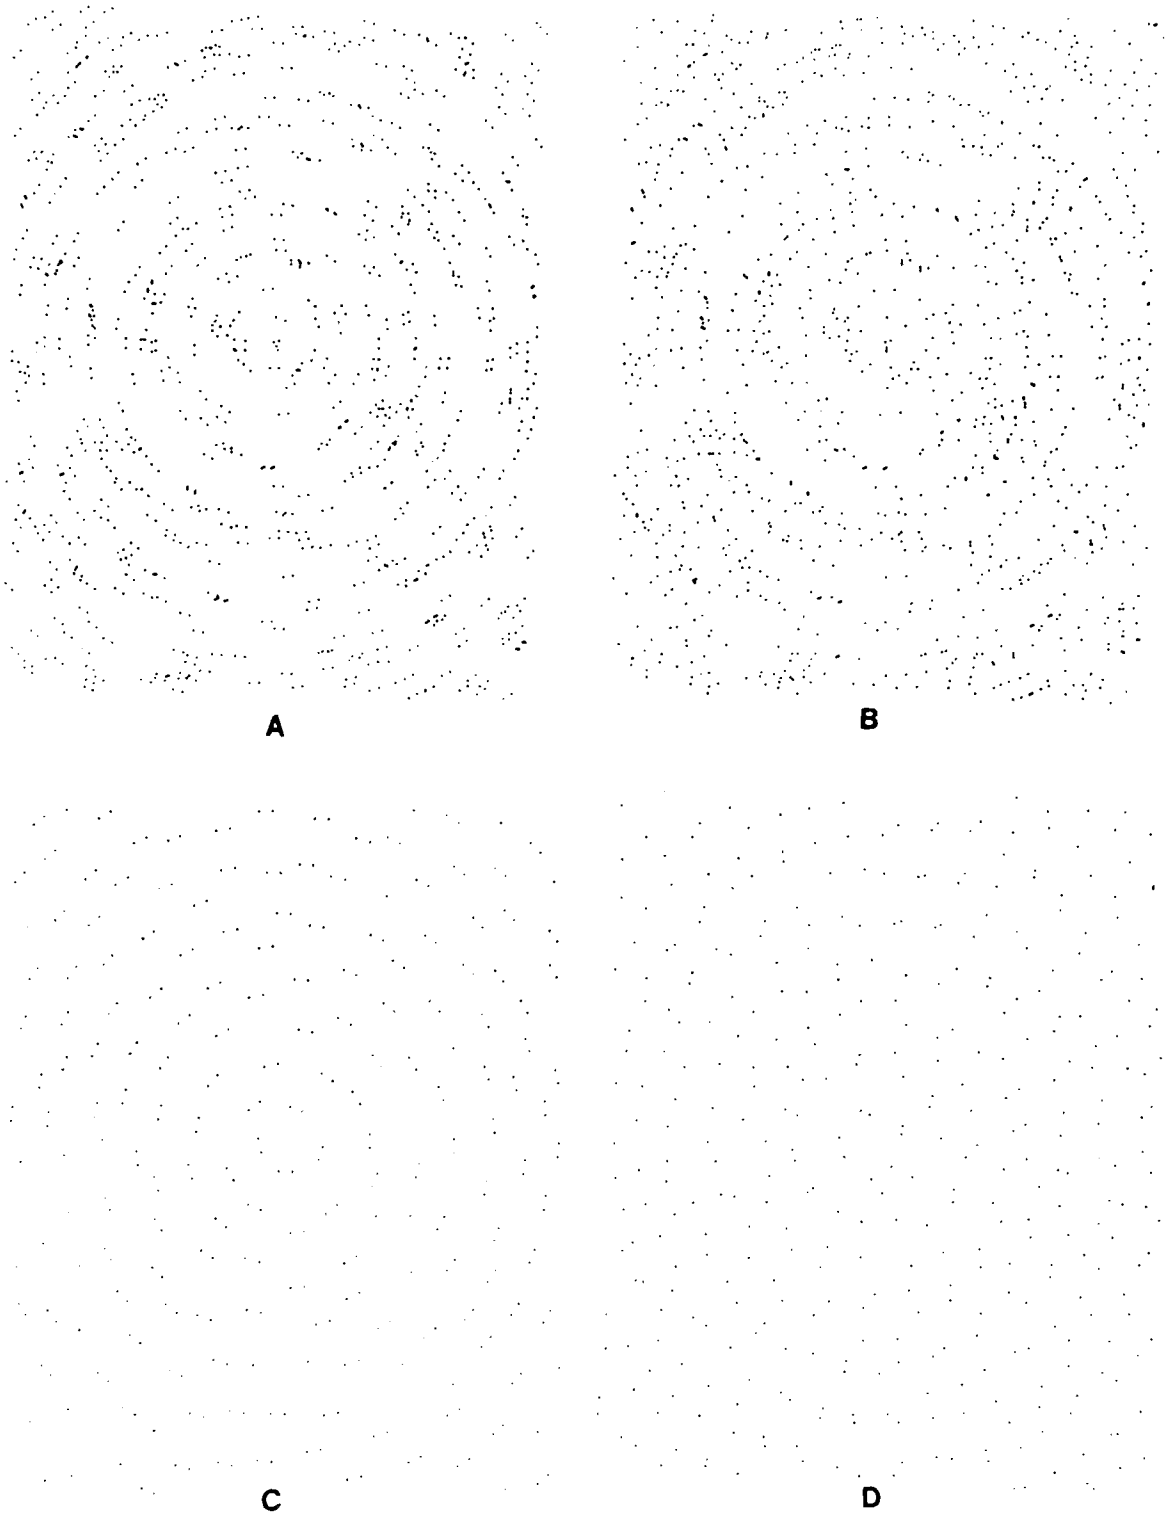

Figure 2. With a random basis pattern, low spatial frequency effects arise even when the dot pairings are no longer evident (compare organization in *a* and *b*). In contrast, *c* and *d* show how the organization in a homogeneous-density Glass pattern is carried by the pairings and not the low spatial frequencies.

3.4.2 Global Organizations may Dominate over Local Pairings

When examining the psychophysics of local dot pairings, it is important to factor out certain global effects that confound the local judgments. In particular, we have found that Glass patterns having foci (such as spiral, radial, and concentric patterns) have a more striking impression of organization than a pure translation pattern. The global impression might well derive from factors other than the apparent local groupings or dot pairings, as just discussed. Hence when using such patterns, the psychophysical judgment of global organization (having subjects distinguish radial versus concentric, for example) might not reveal the local pairings. Since the current concern is how attribute similiarity influences local groupings, we use simple translation patterns for which the global organization is merely parallelism among the individual dot pairs. Later we return to the question of the effect of global organization (see below).

We have also used a pattern, due to Marroquin [1976], which exhibits considerable global organization (figure 3). This pattern has proved useful for examining global grouping tendencies, and in our virtual line modelling, has suggested the presence of long-range processes that detect collinearity among relatively isolated dot pairs. The Marroquin pattern is generated from a square dot grid and two superimposed copies, one rotated 60° and the other 120° relative to the original. All rotations are about the center dot of the first pattern. One may observe in this pattern various types of geometric organization, including circles, rectangles, and more complicated shapes. The pattern also exhibits clusters and voids.

3.5 Evidence for Place Tokens

We will discuss two types of evidence for place tokens. The first concerning the independence

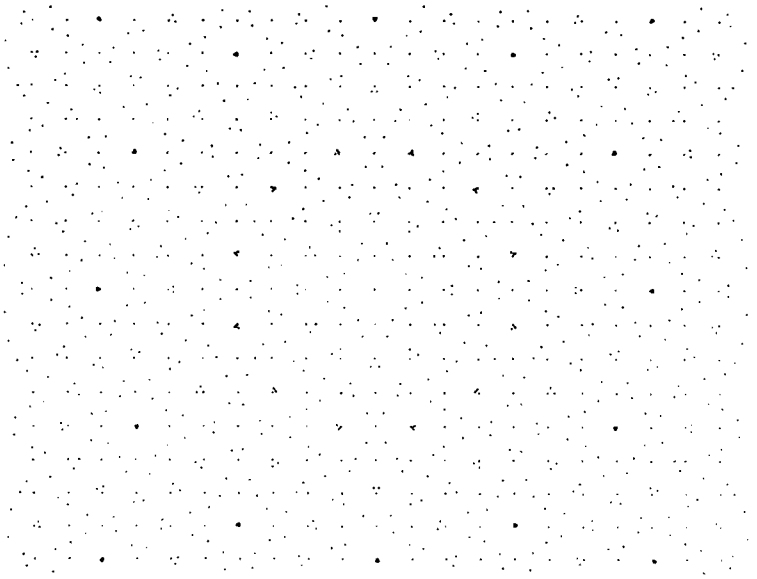

Figure 3. Marroquin pattern.

of the apparent groupings from the spatial frequency content of the patterns. Geometric organization can be seen in patterns that are devoid of low spatial frequencies, which is difficult to attribute to receptive field mechanisms. The second evidence is provided by demonstrating that organization on discrete items is dictated by the similarity of their properties, again contrary to that predicted by linear energy summation models. The apparent independence of grouping processes from spatial frequency content and the importance of attribute similiarity provides evidence for place tokens.

9.5.1 Patterns Devoid of Low Spatial Frequencies

Balanced Checkerboard Dots

It has been shown by Carlson *et al.* [1980] and Janez [1984] that visual groupings can still be seen in high-pass spatial frequency filtered patterns. Their interpretation is that detection of linear features by low spatial frequency tuned channels is not a sufficient explanation for visual groupings, and that some more abstract method is also involved in grouping the discrete items into perceived wholes.

In a similar paradigm, we have examined Glass, Marroquin, and other patterns composed of discrete items that have only very high spatial frequency. The individual texture items are each composed of nine pixels arranged as a 3x3 black-and-white checkerboard with the center pixel the same grey value as the background grey. (Figure 4 shows the configuration). An individual balanced-checkerboard "dot" is intensity-balanced with the background, such that it has subjectively zero average contrast. The background grey is chosen to match the average

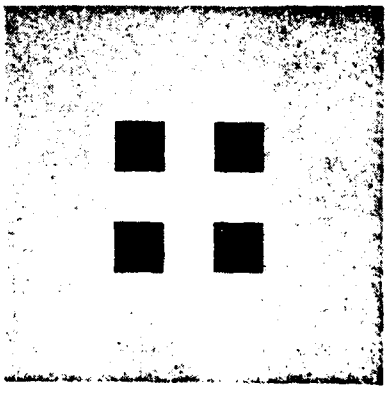

A B

Figure 4. Glass pattern in (*a*) composed of balanced-checkerboard dots (*b*) and thereby devoid of low spatial frequencies.

luminance of the black and white, such that when viewed from sufficient distance the entire pattern appears a featureless grey. For observers that are moderately dark adapted, the individual dots are just visible when they subtend approximately 2.2′ of arc (each pixel subtending roughly .7′ of arc). This scheme provides an advantage over dots that are spatial-frequency filtered or dots with Laplacian or difference-of-Gaussian intensity distributions, in that they present minute but sharp, high contrast detail.

We examined homogeneous-displacement Glass patterns where the underlying dot pattern has constant nearest-neighbor distance thereby providing a pattern of well-spaced, isolated pairs of dots, each separated by a controlled distance (see [Stevens 1978]). When balanced-checkerboard dots are used and the pattern is flashed for 200 msec with masking the local organization is visibly apparent. The dots cannot be resolved beyond a few degrees eccentricity, hence the pattern appears a homogeneous grey except in the vicinity of the direction of gaze, where minute contrast detail is apparent. Because the balanced-checkerboard dots are viewed at the limit of visual resolution, the pattern appears to scintillate in the parafovea. These dots that are clearly visible in the vicinity of gaze appear paired, and the pairs in the vicinity appear parallel. The impression is similar to examining the equivalent Glass pattern composed of regular dots instead of balanced-checkerboard dots.

This pattern in and of itself poses problems for the hypothesis that the dots pairs are detected by simple cells with bar-shaped receptive fields, in the same manner as the demonstrations by Carlson *et al.* [1980] and Janez [1984]. The argument is that the stimuli have insignificant power in the range of spatial frequencies at which a correspondingly scaled simple cell would be expected to respond. More quantitatively, we found that the impression of local organization is apparent when the dots in each pair were separated by at least 30′ of arc. Larger separations of as much as 1° can be tolerated, but the resulting patterns are so sparse that one can resolve only a very few pairs and the impression of preattentive groupings is less compelling, although the task is still performed with short presentation and masking.

Can this very large separation relative to the spatial frequency content be reconciled with the very high spatial frequencies presented by the pattern? If not, it would appear safe to conclude that the perceived pairings between dots arises from mechanisms other than cortical simple cells.

First, consider how a single balanced-checkerboard or DOG dot would stimulate retinal ganglion X-cells of differing central excitatory region diameter ω. An (on-center) X-cell with ω somewhat smaller than an individual pixel would be maximally stimulated when variety located on a white pixel, since the white pixel would fill its excitatory center and black or grey pixels would fall into its inhibitory surround. Larger X-cells would receive progressively weaker stimulation, and those with excitatory centers somewhat larger than an entire balanced-checkerboard dot would produce insignificant response due to the cell's linearity within this range [Enroth-Cugell & Robson 1966; Movshon *et al.* 1978]. This suggests that these dots are detected by receptive fields smaller than Wilson & Bergen's [1979] proposed N channel ($\omega =$

4.4′ of arc), more on the order of the 1.3′ of arc channel proposed by Marr *et al.* [1980]. The 1.3′ of arc channel presumably corresponds to midget ganglion cells, each receiving excitatory input from a single cone.

Consider next the functional organization of a typical oriented simple cell of bar-shaped receptive field (*e.g.* a S_1 cell [Schiller et al. 1976]) with specificity to a bright bar (assume the pattern consists of luminous dots). The elongated excitatory subfield receives on-center LGN X-cell input, while the flanking inhibitory subfields receive off-center input. The critical question is whether, in human vision, one finds bar-shaped receptive fields which summate LGN X-cells having ω approximately 1.3′ of arc over excitatory subfields as long as 30-40′ of arc (*i.e.* over 20 times longer than its width). Negative evidence is provided by the psychophysics of adaptation to gratings and of orientation discrimination in bars. The evidence provided by single cell neurophysiology, discussed below, is more equivocal.

First, in adaptation studies, the summation area over which one finds threshold elevation is spatially limited to an area the size of which is reciprocally related to the spatial frequency, *i.e.* roughly 10 periods in length or width [Howell & Hess 1978; Wright 1982]. The area of functional summation is reciprocally related to the spatial frequency over a large range (4-32 c/deg) of spatial frequencies. This would predict a summation area of roughly 13′ of arc for a channel driven by ganglion cells with ω=1.3′ of arc.

Similar results are found in studies of orientation sensitivity to small bars [Andrews 1967a, 1967b; Vassilev & Penchev 1976; Bacon & King-Smith 1977; Scobey 1982]. They conclude that the receptive fields in human vision that provide information about line orientation have a maximum length of about 9′ of arc in the fovea. These results are directly relevant to the current question, since the stimuli are typically thin bars (2′ of arc width), on the scale of channel that would be responding to the detail within the balanced-checkerboard dots. Burton & Ruddock [1978] have also shown that the threshold elevation effect is length selective when the (bright) bar length is less than roughly three times the bar width. The above studies jointly point to the conclusion that in human vision the receptive fields of the relevant scale are too short to span widely separated checkerboard dots and still be sensitive to the very high spatial frequency content that makes them visible against the background grey.

Turning to neurophysiology, the data are not as conclusive, as one might expect. Consider the dimensions of the receptive fields as mapped conventionally. The minimum width of the central excitatory region is probably the diameter ω of the constituent LGN X-cells [Hubel & Weisel 1962, 1968]. The overall dimensions of simple cell receptive fields, as conventionally mapped in monkey fovea, tend to be somewhat longer than they are wide, from 1/4° by 1/4° to 1/2° by 3/4° [Hubel & Weisel 1962, 1968; De Valois *et al.* 1982]. Within this overall extent, the central excitatory regions are commonly not more than 4ω long; for example, Poggio [1972] reports, for monkey simple cells in the foveal region, receptive field lengths between 6′-24′ of arc (divide by approximately 1.8 for human). On the other hand, Schiller *et al.* [1976, figure 17] show evidence for simple cells that have increasing response as the length of the bar

stimulus increases up to the 6.4 degrees examined (in monkey, at 2-5° eccentricity, therefore one should divide by roughly 4 to compare to human foveal). The evidence from direct receptive field mapping would suggest that foveal simple cell receptive fields of 30′ do not exist in human, but is not conclusive.

Despite the variety of sizes of receptive field that have been mapped in monkey, the psychophysics suggest that in the human visual system the central fovea has a limited range of size of receptive fields that affect spatial frequency adaptation and orientation sensitivity. Hence the balanced-checkerboard dots and the similar demonstrations by Carlson *et al.* [1980] and Janez [1984] show that at least part of the process of detecting collinear (line-like) organization involves explicit groupings on place tokens. It is not enough to detect their organization by spatially blurring the distribution.

Difference-of-Gaussians Dots that Scale with Eccentricity

We next examined whether patterns devoid of low spatial frequencies are effective in producing a global impression of organization. The balanced-checkerboards are so minute as to not be resolved in the parafovea, hence only a small region of the pattern could be discerned at any moment. One could not gain an appreciation for the global organization, of course, given such a restricted effective field of view. Hence we changed from balanced-checkerboard dots to difference-of-Gaussian (DOG) dots, as Carlson *et al.* [1980] used, where the size of the dots scale with eccentricity along the lines used by Wilson and Giese [1977].

We developed the capability to display a dot pattern with DOG dots whose size varied linearly with eccentricity. The pattern would be viewed from a predetermined distance, a fixation point was provided, and the pattern of DOG dots would be shown against a grey background. The scaling function was calibrated such that all DOG dots were equally visible when holding one's gaze at the fixation point, and from a slightly greater viewing distance no DOG dots were visible at any eccentricity.

We found that a Glass pattern of scaled DOG dots presented global organization in much the same manner as a conventional pattern. The radial or concentric arrangement, for example, was apparent in short presentations, even though the individual DOG dots were just discernible against the background grey.

3.5.2 Examining Similarity Grouping

We wish to show that a set of attributes controls the grouping of dots in an image. This set includes color, intensity, size and orientation. We show that each of these attributes is indeed a factor in the grouping process by demonstrating a rivalrous dot pattern, in which only the attribute in question varies between the two patterns, and the visible pattern follows the similarity of that attribute.

After establishing the ability to group dots on each of the attributes we pit each attribute against each other attribute to test their relative strengths. This can be done by embedding two patterns where each is made up of dots having one or the other attribute.

Finally we show the effect of introducing global organization into the rivalrous dot patterns. We use several such organizations, each defined by displacing each dot relative to a focal point. For example, by displacing each dot a fixed angle relative to a focal point a global organization of concentric circles is induced.

Rivalrous Glass Patterns

Our basic paradigm for examining similiarity issues is to generate dot patterns in which two, differently-transformed, copies of a dot pattern are superimposed over the original. Each dot in the original pattern therefore might pair with two corresponding dots; the pattern is rivalrous (figure 5). By displaying the superimposed dots with differing color, intensity and displacements relative to the original dots one can pit, for example, intensity similarity against proximity. This rivalrous pattern was introduced in [Stevens 1978] to demonstrate the basic role of similiarity in the perception of dot pairings, which is being reiterated and extended here.

Because of the strong influence on the apparent local organization induced by certain global organizations we use simple translation patterns. That is, both superimposed copies are simply translated relative to the original by a fixed amount in a fixed direction. Also, we use diagonal translations to avoid the known biases towards the vertical and horizontal. The basis dot patterns have homogeneous dot density, for reasons discussed earlier.

The rivalry between the different organizations in these patterns was tested with a masked tachistoscopic presentation for acurate results.

Proximity and Intensity Similiarity

We established that intensity and proximity differences were sufficient to define pairings in rivalrous dot patterns. In figure 6a the impression is of diagonal lines leading upward to the right — the pairings are between dim dots rather than either combination of bright and dim dots. There is a preference for grouping dots of equal intensity even if there are adjacent dots of greater intensity. This effect, originally reported in [Stevens 1978] is strong evidence against linear-summation receptive field models, which would predict pairings between dim and bright dots, not between dim dots. Prazdny [1984, p. 474] was unable to replicate this result. He provides demonstrations where the global organization (radial versus concentric, for example) does appear to coincide with the dim-bright pairings, from which Prazdny concludes the organization is extracted by operations in the energy domain. Such judgments of global organization, however, are influenced by low spatial frequencies. Recall the low frequency effects demonstrated in figure 2, wherein the overall Gestalt may be apparent even when the local dot pairings are not — the overall effect is likely derived from energy measurements. But since we are concerned with local organization, we use homogeneous-density dot patterns, translation rather

Figure 5. A homogeneous dot pattern with two translated dot pattern results in a rivalrous Glass pattern.

than scaling or rotation, and ask subjects to judge the apparent local organization, not the overall Gestalt, all for reasons discussed earlier. Under these conditions, which we believe more accurately reveal the local phenonema, the pairings are based on similarity measures, not energy (figure 6a).

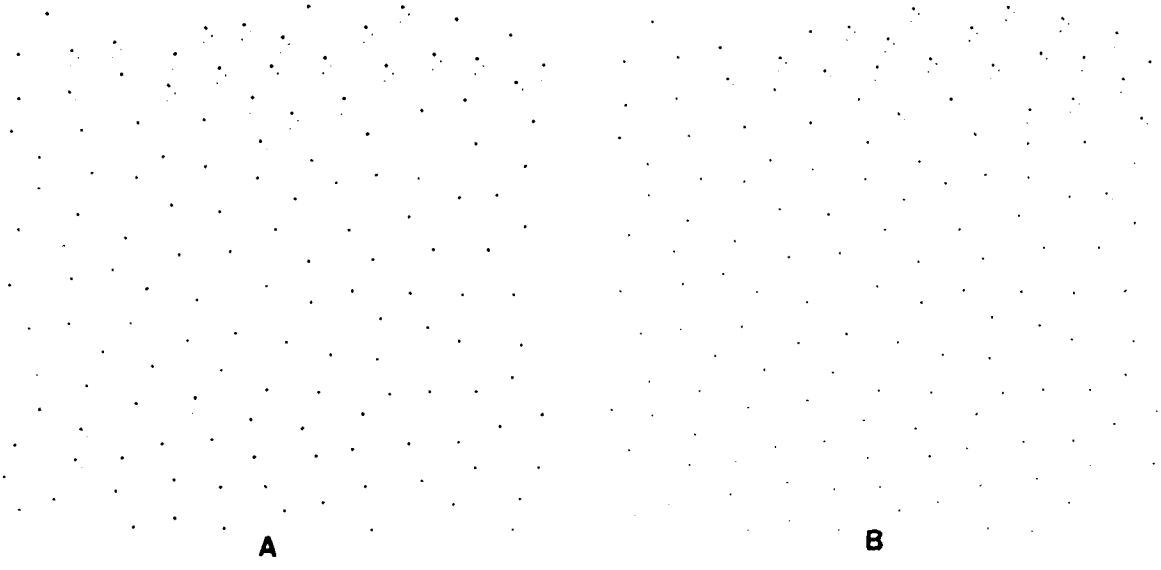

Figure 6. Similar intensity dots are paired, even when dimmer (*a*) and not nearest neighbors (*b*).

There is a well-known preference for pairing dots that are nearest neighbors. The intensity attribute, however, is stronger than that of proximity since in most cases when intensity and proximity are pitted against each other the pairing is between the dots of equal intensity (figure 6b). This pairing can be disrupted if the third dot is much brighter and nearer than the dots of equal intensity. This is possibly due to the apparently global selection of bright dots over dim dots.

Color Similarity

Color similarity appears to be the strongest attribute for pairing place tokens. It is easily established that equal color establishes pairings in the rivalrous patterns. In fact, it is difficult to find attributes that will override the pairing established by equal color. To create a preference for intensity over color the difference in intensity must be great and the pairing must be between the high-intensity dots. We found that there is always a preference for color over proximity which agrees with our earlier findings that proximity was a weak attribute. We also found that we could pit equal color against small differences in intensity proximity and size and still prefer the pairing of the equal-colored dots.

Orientation Similarity

Orientation is another attribute used for grouping. To introduce an orientation into the rivalrous patterns we substituted short narrow bars for the dots. Bars were oriented at 45° and 135°. With other attributes held constant pairing was on the basis of equal orientation. This pairing could be defeated by introducing another attribute such as color or intensity. If the bars of equal orientation were also collinear the pairings were much stronger and required great differences in color or intensity to establish the alternative pairings.

Bars can also pair with dots. In fact the dot-bar pairing can be prefered over the bar-bar or dot-dot pairings. With triples containing two dots and a bar, color, intensity or proximity similarities are enough to establish the dot-bar pairings. With triples containing two bars and a dot the same conditions as above can establish the dot-bar pairings unless the bars are collinear. In this case large color or intensity differences are needed. In the case where the bars are of equal orientation but not collinear the amounts of difference in intensity, color, or proximity must be greater than if the orientations were not the same, but a change in a single attibute is sufficient to change the pairings.

Effects of Global Organization

We mentioned earlier the strong effects introduced when the dot pattern has a global organization. The rivalrous patterns used above had global organization but of a weak sort. The organization was that of parallel lines in a right or left diagonal depending on which dot pairings were seen. The organization was weak and the two possibilities equally likely so that no bias was introduced due to global organization. The following patterns however offer competing global organizations of varying strength. Figure 7a shows a rivalry between a translation and a radial

pattern. One dot in each triple is translated, as in the previous patterns; the other is translated radially away from the center of the image. The displacement is the same in both cases. The dominant impression is of radial organization, even in Figure 7b where intensity similiarity would suggest the translation pattern (as was the case in figure 6).

A stronger global organization than the radial pattern is that of concentric curves around some focal point. Figure 8 shows rivalrous patterns where the translations are radially away from the center and through an arc relative to the center. The impression is of concentric circles in all cases. The pairings that result in the concentric circles resist changes in intensity, proximity and color. Even when the initial impression is of radial lines, that is when there is

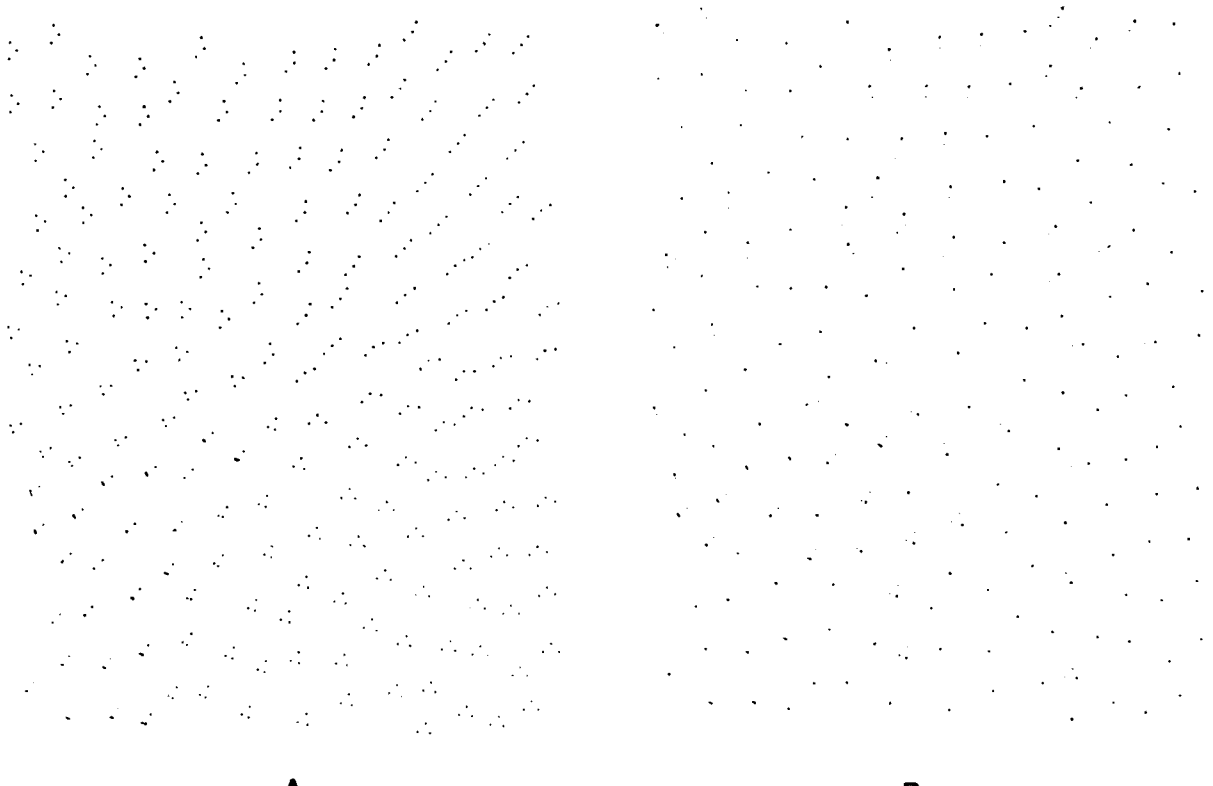

Figure 7. Examples of rivalry between pure translation and radial patterns.

enough similarity in the radial direction, one can attend to the concentric circles and reverse the pairings.

An even more striking rivalry is made by using triples of points forming an equilateral triangle with one side of the triangle oriented in either the radial or the concentric direction. Figure 9 shows this rivalry with the triples oriented in the radial direction. In figure 9a the dots forming the radial lines are brighter than the remaining dot. The radial pattern is visible although the competing spiral lines are also visible. In figure 9b all the dots are the same intensity and the radial line are virtually impossible to see.

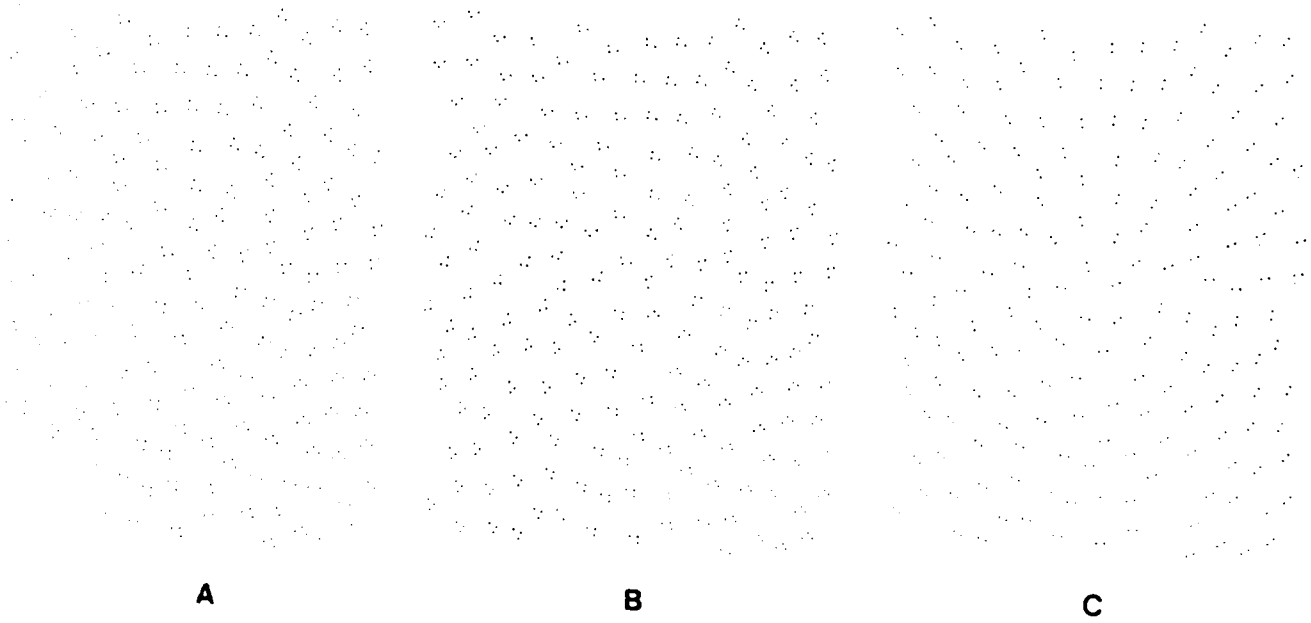

Figure 8. Examples of rivalry between radial patterns and concentric patterns.

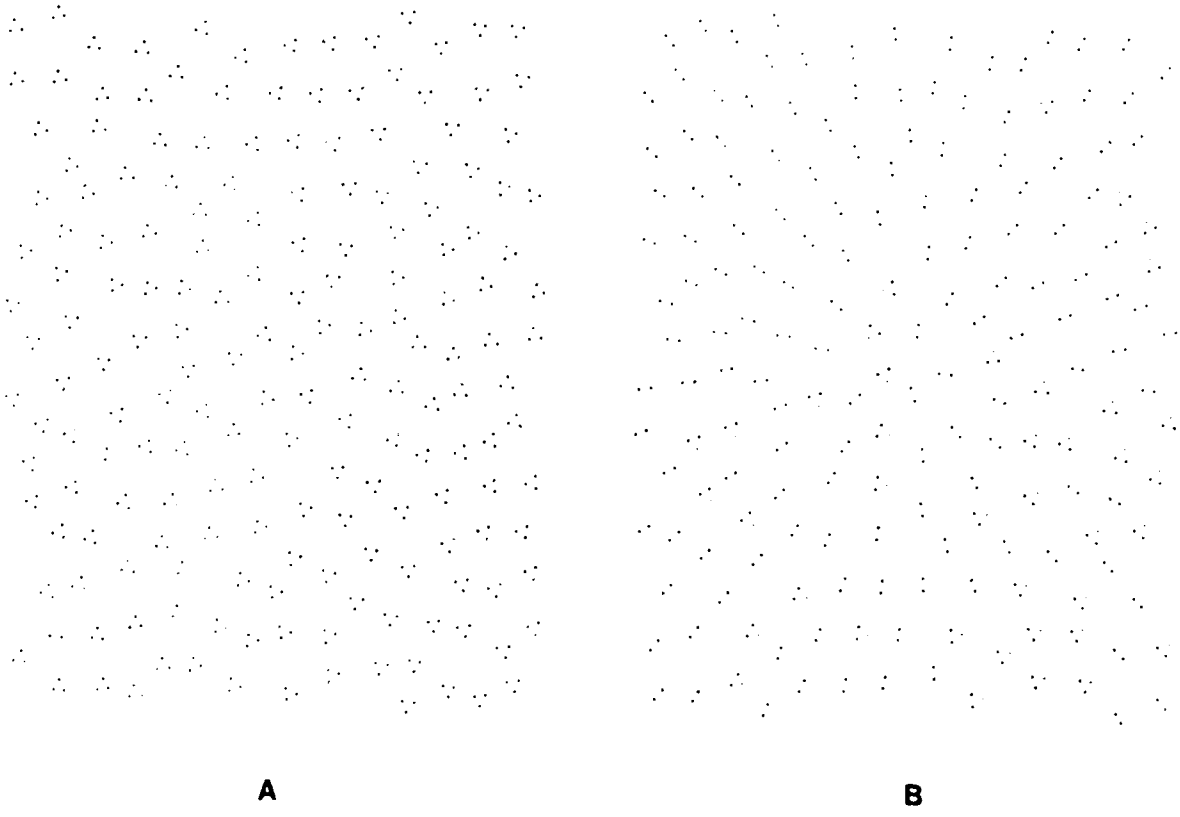

Figure 9. Examples of rivalry between radial patterns and spiral patterns.

3.6 Extracting Structure

3.6.1 Attribute selection

One of the principal computational problems of the primal sketch is to construct descriptions of global organization on the basis of local evidence. The task must be achieved, at least in part, by "bootstrapping", as it were. An issue is how to initiate this process. Marr [1982, p. 47] observed that

> The items generated on a given surface by a reflectance-generating processing acting at a given scale tend to be more similar to one another in their size, local contrast, color, and spatial organization than to other items on that surface.

As a strategy for detecting structure, we suggest that this observation can be inverted in the following way. Similarity and structure are generally correlated, i.e.

> The items at a given scale that are similar in size, local contrast and color tend to have a common spatial organization.

That is, (geometrically) structured intensity changes are probably similar along various (non-geometric) attribute dimensions, while uncorrelated intensity changes, such as those arising from different physical causes, are not expected to be similar except by coincidence. Hence:

> Similar intensity changes in any locality are likely to be structured as well. A good place to look for geometric organization is in those subpopulations defined by similarity.

Thus as an early step in the bootstrapping process of extracting structure we suggest the visual system seeks evidence of similarity in various non-geometric dimensions to partition the intensity changes into subpopulations which are subsequently analyzed for geometric organization. And the more similarity shared by a set of elements (edge segments, say) the more likely they are physically correlated, clearly.

Evidence for similarity is therefore sought by examining the spatial distribution of values of certain attributes (such as color, contrast, and so forth). A prominent band or peak within the distribution is likely to reflect related items. In terms of feature maps, wherein properties are mapped out in distinct representations, one might expect analysis of individual maps, with perhaps mutual support across maps. Recalling the earlier discussion of selection, by this proposal a distinguishable "signal" in any selectable attribute may initiate the bootstrapping.

9.6.2 Virtual lines represent pairings of selected items

Having selected a sub-population of elements of similar attribute, individually represented as place tokens, their local arrangement must be represented. Virtual lines could make explicit a pair of locations, an angle, and potentially a length. Local operations on virtual lines would then begin to build geometric relations among tokens. Parallelism, as among dot pairs in the Glass pattern, could be detected by selecting those virtual lines that share the same orientation as the prominent virtual line orientation within a given spatial region [Stevens 1978]. Note the repetition of the general theme of computation, selection, computation, ... Likewise, collinearity, such as among the dots in the chains also observed in Glass patterns, could be detected by pairwise collinearity of the corresponding virtual lines.

Virtual lines appear to be piecewise-linear constructions spanning the selected place tokens, not continuous curves (e.g. splines). This observation poses a particularly challenging issue for algorithms that construct virtual lines, because the connections or lines that are to be constructed need to connect those neighboring tokens that are collinear, but not necessarily isolated, and not necessarily at any given, fixed, separation. We have performed computational experiments to determined the utility of constructing virtual lines by spatial blurring of place tokens (n.b. we are not referring to low spatial frequencies in the image intensity domain, but

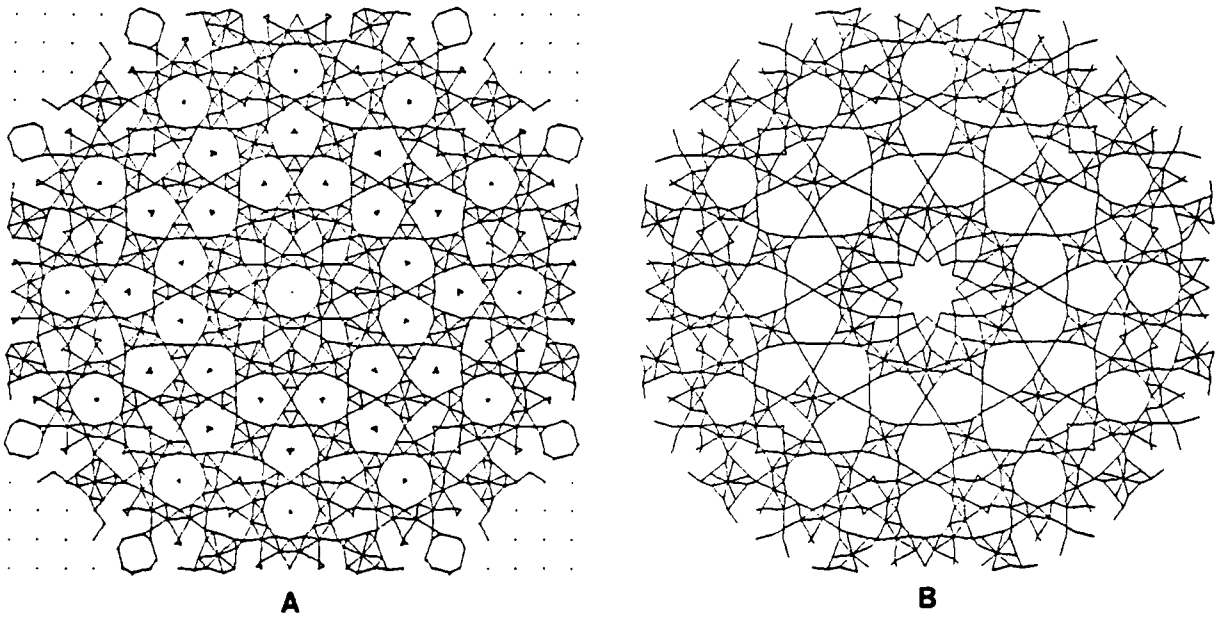

Figure 10. Of the virtual lines constructed on a Marroquin pattern in *a*, those that are pairwise collinear are selected in *b*.

in a blurring of the discrete token array). We have found that blurring schemes are successful only for isolated pairs or chains of dots. Collinear dots imbedded in a background of extraneous dots, as are the circles and squares observed in the Marroquin pattern, are not successfully extracted by such means. Spatial blurring (whether in the image intensity or place token domain) only serves to find density inhomogenieties, such as clusters and voids.

From exploration of dot pairing algorithms, we conclude that the method for constructing pairings in human vision likely has the following properties: an isotropic (circular support) for detecting nearest neighbor to a given token, a proximity measure that scales (implicitly or explicitly) with the density of selected tokens. (Recall that brighter dots can be selected independently of the number of nearer dimmer neighbors, suggesting that the groupings among these dots occur after selection.) While dot pairings are substantially independent of scale, "relatively isolated" dots (compared to the mean dot density) are not paired with neighbors.

In [Stevens 1978] a simple virtual line algorithm was proposed in order to show that iterative, relaxation or cooperative processes are not needed to extract local parallelism, as exemplified by the Glass patterns. A local process first defines virtual lines between neighboring tokens, along the lines of [O'Callaghan 1974a, 1974b, 1975] where all neighbors within some factor k time the nearest neighbor distance are connected. Then the virtual lines are histogrammed in each vicinity, resulting in a peak orientation if parallelism is present. It is then a simple matter to select those virtual lines that have approximately the same orientation as the

local peak. The algorithm was not posed as a model for the detection of local parallelism in these patterns, but rather as a demonstration that a simple, noniterative computations might serve (see [Marr 1982] for discussion).

This virtual line algorithm was applied to the Marroquin pattern (figure 10a), and in a second step, those virtual lines that are pairwise collinear to within $\pm$ 20-30° are selected (figure 10b). The algorithm is rather successful in making explicit the pairwise connectivity among dots seen in the Marroquin pattern. Observe that the local connectivity is established by this algorithm, and it requires only extraction of extended chains of virtual lines to make explicit the various curvilinear and rectilinear figures seen in this pattern. The algorithm has also applied to various random patterns including the experimental stimulus patterns Caelli *et al.* [1978; figure 4] has published (figure 11a). In figure 11b the algorithm constructs virtual lines and draws with solid lines those that are pairwise-collinear

The conclusion, thus far, is that approximately collinear, neighboring tokens are visually quite salient, and these local organizations appear amenable to parallel computation with a local support. There is considerable work needed to determine what aspects of the algorithm contribute this apparent success, and to justify what is otherwise an *ad hoc* algorithm.

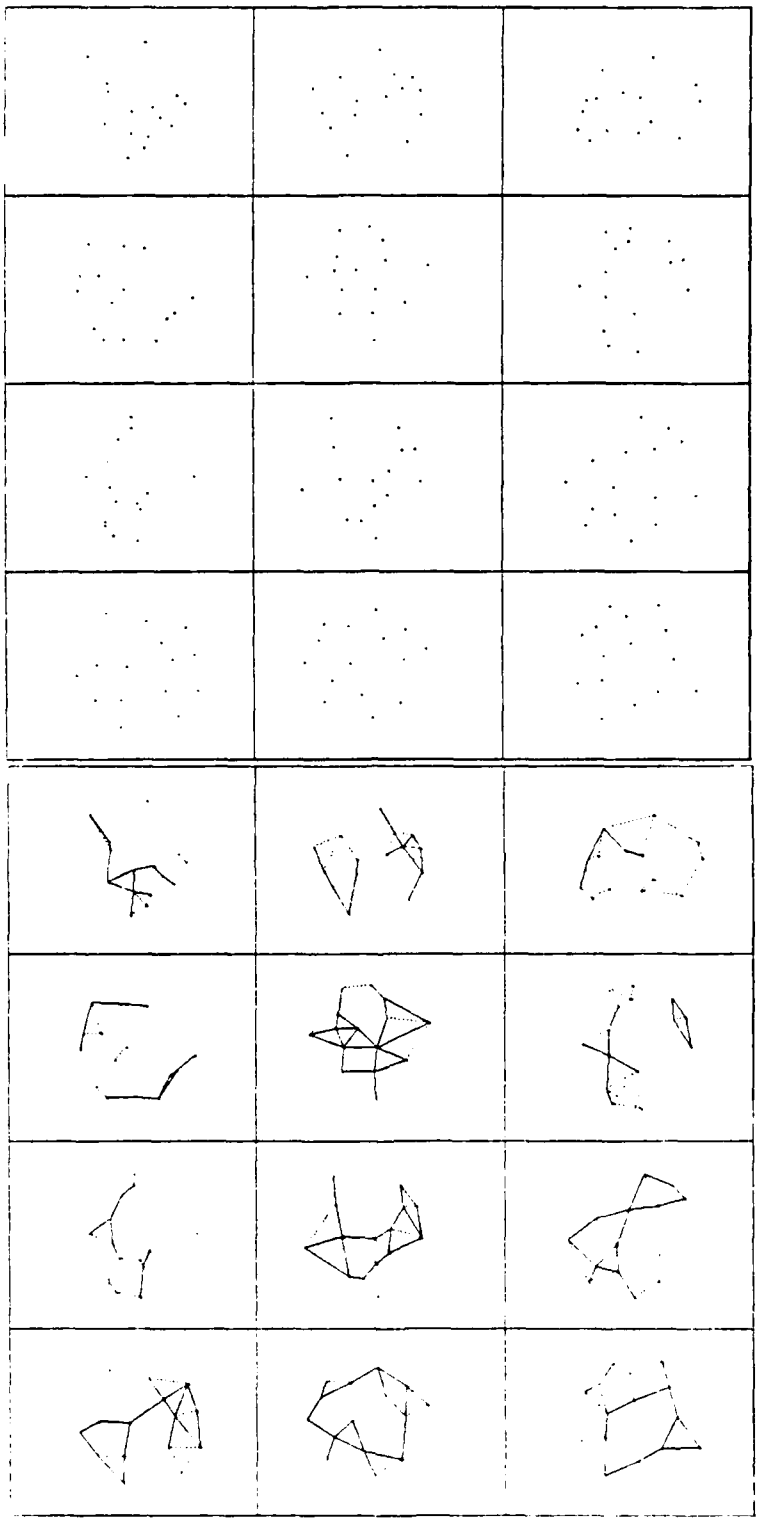

Figure 11. Dot grouping stimulus patterns in *a* from [Caelli *et al.* 1978, figure 4]. In *b* the pairwise-collinear virtual lines are displayed (see *text*).

4. THE CONCAVE CUSP AS A DETERMINER OF FIGURE-GROUND

4.1 Introduction

Figure-ground, the process of distinguishing a figure relative to its surround, is often exemplified by Rubin's [1958] vase-face illusion. The contours which, in one figure-ground interpretation, comprise the silhouette of a vase, can also be seen as the edges of two faces in silhouette. The figure-ground interpretation in such illustrations is readily affected by focussed attention and verbal suggestion. Figure-ground presumably has a substantial preattentive component as well, with factors such as size, brightness, symmetry, closure, regularity, and convexity contributing to an "immediate" impression of figure against ground. Random texture tends to segregate perceptually on the basis of contrast, such that the white patches might appear as figure against the grey and black patches, or the black patches against the white and grey [Julesz 1965; Richards & Purks 1978]. For such random texture there is a tendency to see white-as-figure and to see smaller-as-figure [Rubin 1958]. These tendencies are balanced when roughly 0.4 of the area is black [Frisch & Julesz 1966]. Symmetry, if present in the texture, dominates size (and contrast), causing one to prefer the symmetric forms even if larger (or darker) [Hochberg 1964]. Still more important is convexity, which is favored over both regularity and symmetry [Kaniza & Gerbino 1976].

The importance of convexity in determining figure-ground is demonstrated in figure 1, generated by placing convex blobs of various shapes at random locations. The immediate interpretation is usually of a texture comprised of small convex objects, regardless of the contrast sign (compare figures 1a and 1b). Note that virtually all of the texture area (81%) is seen as figure; only the small irregular fragments are seen as ground, in strong distinction to the trend reported by Frisch and Julesz [1966]. Of course, while the placement of the blobs is random in figure 1, the resulting texture is not random, owing to the regular, convex, geometry of the individual blobs. For comparison, the texture in figure 2 is composed of irregular, concave, blobs. Here the figure-ground sense is more ambiguous and exhibits the tendencies Frisch and Julesz [1966] discussed. It is easily shown that convexity becomes important only when the two figure-ground alternatives are spatially adjacent, so that the assignment of contour to figure is locally ambiguous. When the same elements as comprise figure 2 are sufficiently isolated they can be seen as figure even though they are concave.

Convexity, as a determiner of figure-ground, derives from the fact that most physical textures are composed of compact surfaces that are convex almost everywhere. They give rise to image texture that, in general, have convex outlines or silhouette contours. But how is convexity measured and processed by the visual system?

4.2 Computational Issues

4.2.1 The Texture Parsing Problem

Most textures are the image projections of discrete and compact physical objects. An image of

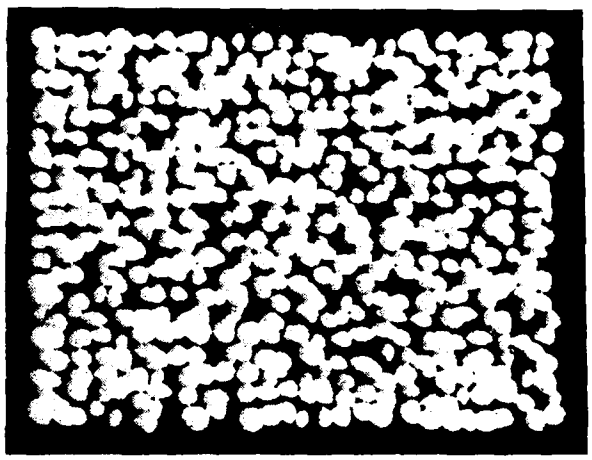

A

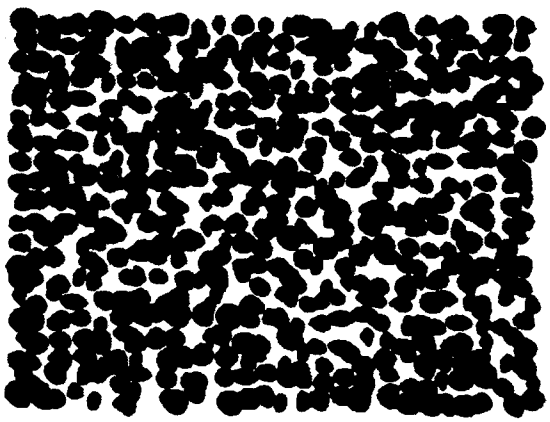

B

Figure 1. Figure 1. Texture perceived as small convex blobs in close packing, independent of contrast sign.

---

leaves against the bright sky, for instance, results in a mottled texture of leaves silhouetted against a bright background. In the early visual processing of such an image, there is presumably a representation of the intensity changes corresponding to the leaf silhouettes. Locally, a binary decision is necessary as to which side of the edge corresponds to figure. This decision effectively parses, two-dimensionally, the set of edges into one of two possible figure-ground interpretations. For the silhouetted leaves, the correct parsing treats the dark regions as figure; the alternative (and incorrect) parsing treats as figure the fragments of sky visible between the leaves. The texture elements that result from the latter parsing, of course, do not correspond

to individual physical objects; the shape of each is generated by the placement, shape, and orientation of the surrounding leaves.

The shape properties of the visible portion of the silhouette depend on which side is seen as figure [Hoffman & Richards 1982, 1984], and thus the visual description of texture depends on the figure-ground interpetation. Moreover, geometric shape properties are meaningfully attributed only to the images of physical objects (e.g. the leaves and not the random shapes of the interstices). Therefore, visual processes that take as input a geometrical description of image texture (such as processes that lead to texture segmentation and recognition) depend on

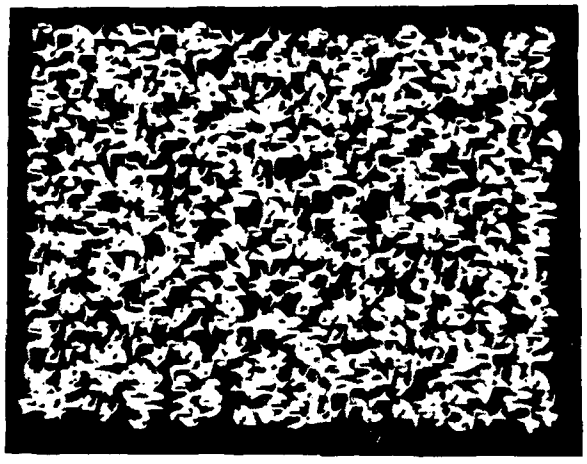

A

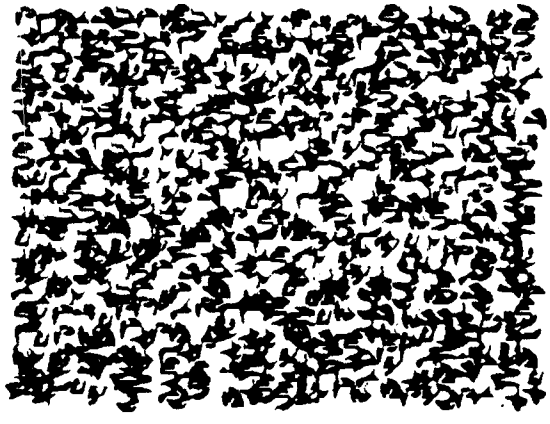

B

Figure 2. Texture composed of concave, randomly shaped blobs in close packing. Observe the figure-ground dependence on contrast sign (see text).

first achieving the correct figure-ground parsing. The visual system has apparently developed robust strategies for deriving texture descriptions which reflect the order, regularity and structure of the objects comprising the physical texture, and which ignore the randomness of the interstices. Note that two explanations can be forwarded: either the visual system derives two texture descriptions in each locality (corresponding to the two texture parsings) and selects that which reflects the structure and recognizable shapes of the physical texture, or it determines the correct parsing in a primarily "bottom-up" manner.

The consequences of figure-ground ambiguity in texture is not widely recognized. Two probable reasons are that we are very adept at deciding figure-ground and seldom see reversals of figure-ground in natural scenes, and for synthetic textures in particular, the discrete constituent elements (the bars, dots, etc.) are usually sufficiently separated that the figure-ground distinction is again unambiguous and stable. It is in consideration of natural images that the significance of the computational problem becomes apparent.

4.2.2 Computing Convexity in Texture

The computational problems associated with figure-ground raises substantial questions regarding how convexity might actually be measured. Convexity is a global property of a closed curve, which has several mathematical definitions, e.g. *i)* a straight line connecting any two points on the curve lies entirely within the curve, *ii)* the curve has no inflection points, and equivalently, *iii)* the curvature has everywhere the same sign. Note that if (i) is rephrased that a straight line connecting any two points on the curve does not intersect the curve at any intermediate point, then all three definitions can be applied to open as well as closed curves. Each of the above definitions might suggest a variety of algorithms for determining convexity.

Consider a smooth arc of curve without inflection points. Assuming it corresponds to the physical edge of a convex object, it immediately follows that the convex side of the curve is figure. The evidence provided by curvature sign alone is very local, as if the curve were examined through a small aperature. It is akin to deciding figure-ground on the basis of contrast sign, where presumably the lighter side of an edge is more likely to correspond to a physical object (and the darker to be shadow or background). Just as one can demonstrate the tendency to interpret lighter as figure one can demonstrate the tendency to interpret as figure the convex side of a curve (see below). Since this tendency persists in stimuli that are devoid of other relevant information, one may conclude that some measure of curvature sign is a salient property as regards figure-ground. But we will also show that curvature sign is not the only factor underlying the convexity preference.

The preference for convexity likely stems from several causes, all of which are consequences of convex objects imaging as convex silhouettes. In addition to the straightforward notion of curvature sign just discussed, there are localized events that arise where convex silhouettes overlap which also indicate the appropriate figure-ground assignment. Since the physical objects that comprise a texture are usually distributed three-dimensionally in space,

their silhouettes often overlap in the image. At each point of overlap the two silhouettes conjoin to produce a sharply discontinuous concave cusp. The figure-ground interpretation of a concave cusp is straightforward: the region within the cusp is ground, and the two component arcs correspond to two physically distinct objects. The cusp point itself has no physical significance, however, it is merely the point where they overlap from the given perspective.

The figure-ground stimuli that Kaniza and Gerbino [1976] used to show the preference for convexity have, in addition to smoothly convex curve arcs that define and enclose convex shapes cusp-like discontinuities in the curves. We suggest that these sharp cusps contribute strongly to figure-ground. The observed convexity preference is due not only to curvature sign along smooth arcs of figures, but to the geometry of the sharp discontinuities where the smooth arcs conjoin.

4.2.3 Figure-ground Determination and Part Boundaries

The geometry of the concave cusp suggests the figure-ground relationship across each of the two arcs. It also demarks a point where two distinct silhouettes intersect, the physical interpretation of which is that distinct physical parts, either separated in space or abutting, project so that their silhouette contours intersect. Not only are the two figures distinguished from their common background, but from each other. This latter role of the concave cusp, as *a priori* evidence for a part boundary, has been recognized by other researchers, but for a distinctly different physical interpretation. Specifically, we are concerned with the cusp as evidence that two convex objects partly overlap or abut; the other work has concerned parts that interpenetrate or join to form a common object and moreover the parts need not be convex.

The silhouette of an object composed of distinct parts generally carries information about where the parts conjoin. Deep concavities in the silhouette outline are good candidate part boundaries, specifically at the point where the curvature is most negative at the base of the concavity [Marr & Nishihara 1978]. Hoffman and Richards [1984] similarly observe that points of minimum curvature along a silhouette curve often correspond to part boundaries, citing the fact that for two arbitrarily shaped surfaces that interpenetrate, the locus of intersection is a contour of concave discontinuity of their tangent planes. Consequently they propose parsing surfaces in 3-D along loci of negative minima of each principal curvature, and in 2-D, at points of negative minima of curvature [Hoffman & Richards 1982, 1984]. Their treatment of 2-D silhouette contours is derived from the 3-D case. Their proposal results in apparently psychologically valid predictions both for the apparent parts of surfaces in 3-D and of curves in 2-D. They define curvature sign relative to the side of the curve that is regarded as figure (a protrusion in the silhouette is associated with positive curvature by their convention, and an indentation with negative curvature). If the figure-ground sense across the contour reverses, the curvature sign also reverses (so that minima of curvature become maxima of curvature and vice versa). Since what is regarded as a minimum of curvature depends on the figure-ground interpretation, they predict that the subjective parts of a curve will appear delimited by minima of curvature as

dependent on the figure-ground sense, a prediction that is upheld in e.g. the vase-face illusion.

The interpretation of a silhouette curve thus depends on the figure-ground and assumes that figure has been distinguished from the background. We propose the following extension: that certain types of cusp are early *a priori* evidence for which side is figure, as a precursor to subsequent analysis of the silhouette's shape.

4.2.4 Figure-ground Interpretation of Cusps of Different Type

Not all cusp discontinuities in tangent along a curve can be interpreted equally as evidence for figure-ground. Even the concave cusp as we define it has an alternative interpretation, that being of a sharp concave physical object such as a thorn (see figure 3a) rather than a gap between two convex objects. That either interpretation might be valid cannot be ignored, but in the absence of other figure-ground evidence, a concave cusp more likely corresponds to two overlapping or abutting convex objects than to a single sharp concave spike or thorn. By this argument, the observed bias or preference for convexity reflects a statistical fact about our visual world.

The concave cusp, formed by two convex silhouettes, is but one of six types of cusp, the geometry of each depending on the combination of curvatures of the two intersecting silhouette curves at their point of intersection. Since each arc may have either positive, negative, or negligible (zero) curvature, the concave cusp is termed negative/negative and the other cusps are: negative/zero, negative/positive, positive/zero, positive/positive, and zero/zero (see figure 3). These five cusp types are weaker figure-ground evidence than the negative/negative, and can be rank-ordered roughly. The negative/zero cusp is somewhat weaker, as only one convex object would be involved, and the straight edge providing no additional information. Still weaker evidence would be the negative/positive cusp which, by the above interpretation would correspond to the intersection of a concave and a convex silhouette. The more probable

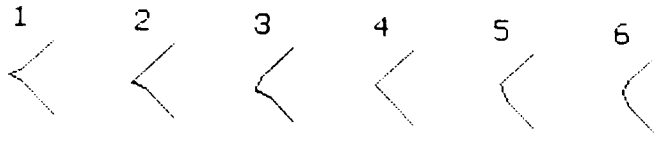

Figure 3. Six types of cusps: *1*) negative/negative curvature, *2*) negative/zero curvature, *3*) negative/positive curvature, *4*) zero/zero curvature, *5*) zero/positive curvature, *6*) positive/positive curvature.

interpretation of the negative/positive cusp seemingly would be a sharply pointed figure such as a leaf tip. Similarly, the positive/positive and positive/zero cusps are less unlikely to be part boundaries than sharp convexities in the silhouette of a single figure. Finally, the zero/zero discontinuity, a simple corner defined by two straight edges, is clearly the weakest evidence.

In the following we show that the known figure-ground preference for convexity (e.g. reported by Kaniza and Gerbino [1976]) derives in part from the presence of concave (*negative/negative*) cusps. Since the concave cusp is comprised of two convex arcs, and we recognize that the convexity of an individual arc induces a figure-ground preference in and of itself, we must, in the process, show that the particular geometric arrangement of cusp, and not merely the curvature of the arcs, is effective.

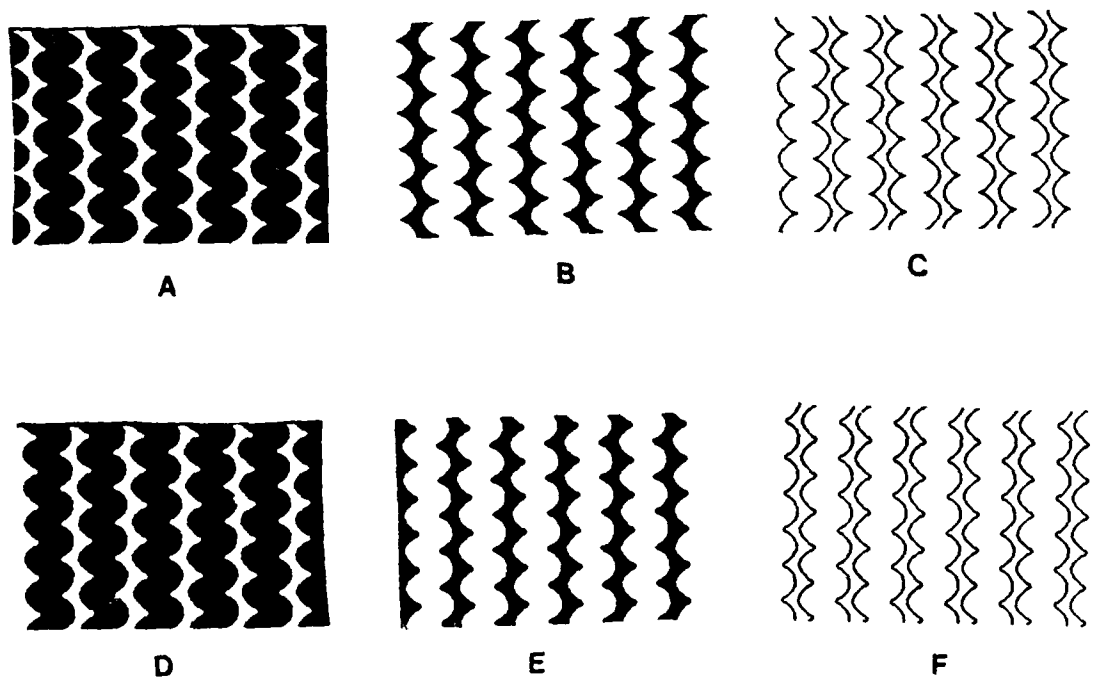

Figure 4. Asymmetric convex (sinuous) shapes vs. asymmetric concave shapes. In *a,b,c* there is preference for the sinusous shapes due to the sharp cusps. In *d,e,f* there is still a preference for sinuous shapes but the alternative parsing can also be seen, especially in *e*.

4.3 Demonstrations

Figure 4a presents a texture with two distinct figure-ground parsings. When white is regarded as figure one sees sinuous shapes resembling coiled telephone cords; when black is figure one sees concave shapes such as thorns. There is a preference to see the sinuous shapes, wherein the concave cusps demark convex segments along the cords. The contrast is reversed in figure 4b, and as expected, the concave thorns, now white, may be seen as figure more readily. The competing contribution of contrast sign is removed in the line-drawn version in figure 4c and the sinuous shapes are again dominant. Neither figure-ground interpretation is absolute in figure 4a-c, however. Spontaneous reversals are frequent, nonetheless the initial impression is generally to see the sinuous, convex shapes, and this interpretation is held the greater fraction of the time. In figure 4d-f the pattern has been subtly modified to remove the sharp concave discontinuities; the figure-ground interpretation is now more ambiguous. Compare the line drawings in figures 3c and 4f to observe the effect of the sharp cusps. We suggest that the relatively greater stability of the sinuous (coiled telephone cord) interpretation in figure 4c over figure 4f to the presence of sharp the concave cusps in figure 3c and their absence in figure 4f — the change in terms of total curvature being negligible, and other factors remaining constant.

Convexity and symmetry can be placed in opposition to further increase the figure-ground ambiguity. In figure 5a, for example, one may see white sinuous cord-like figures, as before, where the convexity dominates. But observe that the black background shapes are symmetric, and if seen as figure resemble strings of concave beads. The concave but symmetric figures are more readily seen when the contrast is reversed as in figure 5b; the line-drawn version is shown in figure 5c. In figure 6 both figure-ground interpretations are symmetric, and the tendency is to see convex beads rather than the alternative concave beads. The convex figure interpretation is strong even in figure 6d where the spacing favors the alternative concave-figure interpretation. There is still a substantial tendency to see strings of convex beads despite the wide separation of the convex arcs that define their silhouettes. In figure 6e, where the sharp con-

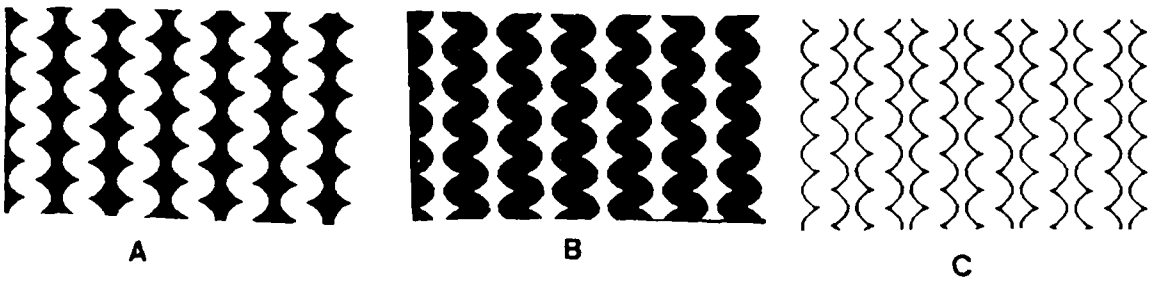

Figure 5. See text.

cave cusps are slightly rounded, the alternative interpretation of widely separated strings of concave beads is more readily achieved.

While an effect due to the sharp concave cusp is evident in these demonstrations, the figure-ground interpretations are influenced by many uncontrolled factors, particularly since the resulting shapes are two-dimensional. To further simplify the stimulus, therefore, we next turn to one-dimensional patterns (see figures 7 and 8). These patterns, being line-drawings, remove the influence of edge contrast on figure-ground, but are interpreted in terms of occluding edges nonetheless. Moreover, rather than defining competing two-dimensional shapes, the patterns of curves merely suggest a one-dimensional arrangement of overlapping edges. The local figure-ground problem is thus reduced to determining whether the curve corresponds to a left or right edge.

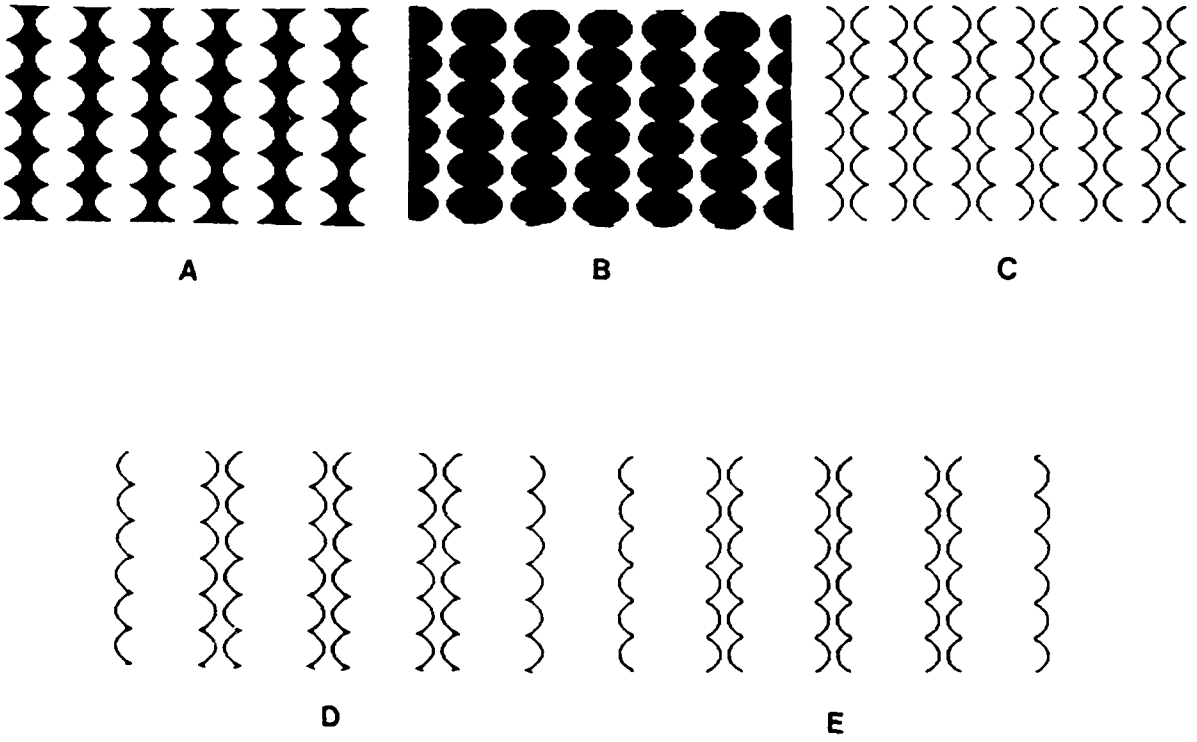

Figure 6. See text.

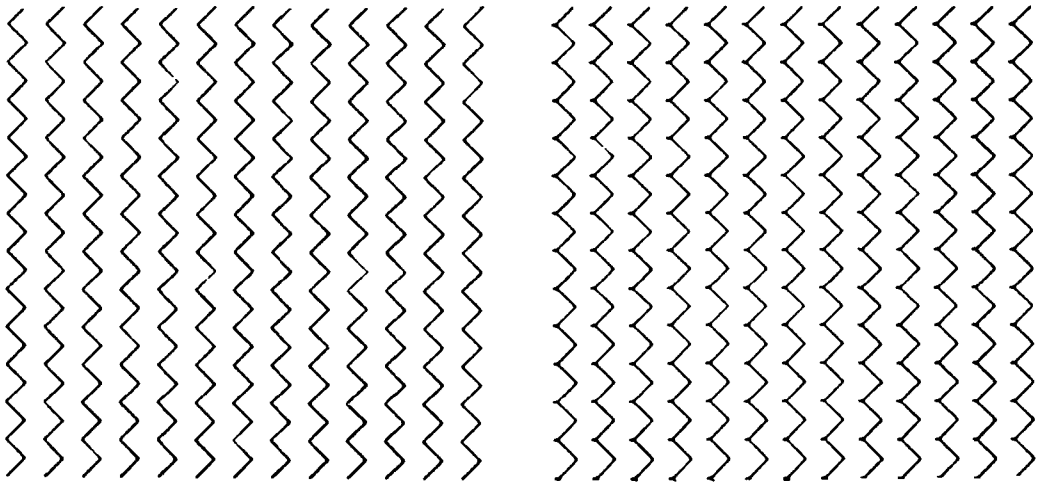

Figure 7. The symmetrical pattern in *a* can be seen as a series of serrated edges that successively overlap. There is roughly equal preference for overlap to the left as to the right. The slight cusp introduced in *b* causes strong preference for edges to overlap to the right.

Consider first the symmetrical pattern in figure 7a, which can be seen as serrated edges (sawteeth) overlapping either to the left or right with equal ease. In figure 7b a very slight concave cusp is introduced by appending a minute line segment on the left-hand vertices, with the result that one sees an arrangement of occluding edges, each partly overlapping the surface below it to the right and in turn occluded by the nearer surface to the left. (This will be termed "overlapping to the right". The left side of each curve is seen as figure against the background immediately to its right.[1]

The geometric arrangement in figure 7b, which influences apparent figure-ground so strongly, is composed of straight line segments instead of continuously curved arcs (as in figure 4a). This arrangement serves us in two ways. First, it permits patterns to present distinct and effective cusp-like features while controlling for contour curvature. Second, they provide insight into the defining geometry of the cusp feature, as will be discussed later.

It was mentioned that since a concave cusp is formed by the intersection of two convex (positive curvature) arcs, we need to show that contour curvature alone is not determining figure-ground. This was demonstrated, in part, by figures 4c versus 4f where blunting the concave cusp adds little to the convexity but extinguishes the discontinuity, and hence, we argue, the concave cusp feature. Also consider figure 8. In figure 8b one prefers overlap to the left

[1] Note that if the patterns were oriented horizontally the judgments of overlap direction would be confounded by the independent tendency to interpret depth as increasing as one scans from bottom to top. The apparent overlap is biased towards interpreting the each edge as occluding that which lies above it (rotate figure 7 so that the sawtooth curves are horizontal).

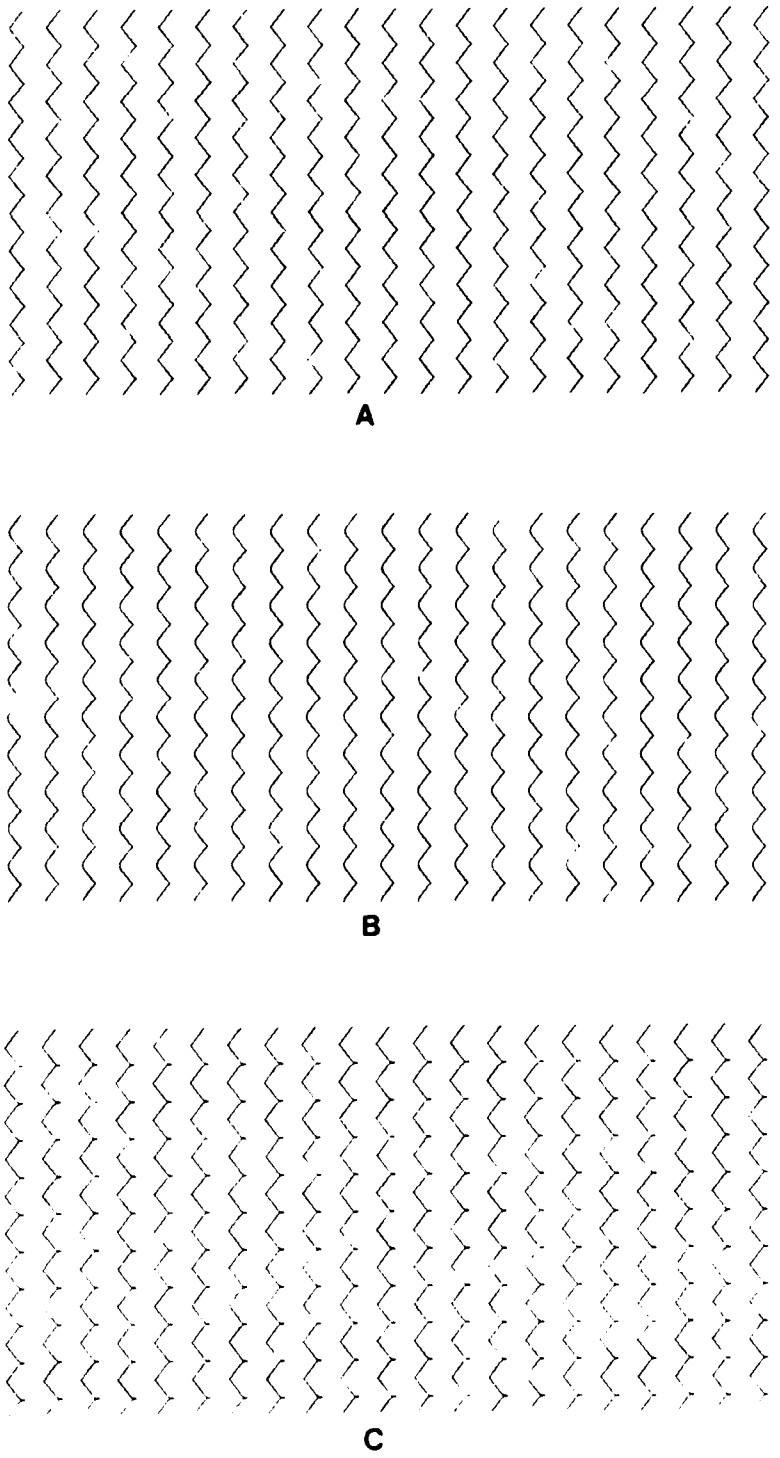

Figure 8. The symmetrical pattern in *a* is ambiguous. The shallow convexities in *b* induce a preference for overlap to the left. The primitive concave cusps in *c* have similar effect, and are more effective in short presentations.

on the basis of convexity; note that only a slight amount of curvature is needed to produce this bias (compare with the symmetrical sawtooth pattern in figure 8a). In figure 8c the overlap is at least as strongly biased towards the left by the straight-line approximations to concave cusps. These cusp shapes induce a strong impression of figure-ground in the absence of smooth convexity. The following reports on experiments that further explores these effects.

4.4 Experiments

Two experiments were performed. The first involved subjects judging the direction of overlap, as in figures 7 and 8 above. From the direction of overlap one can infer which side of the curve was regarded as figure. It can also be determined by presenting a probe dot on one side or the other of a given curve within the stimulus figure and, since the curve is interpreted as an occluding edge, have the subject respond whether or not the dot was on the edge (figure) or on the background. The second experiment employed this task.

4.4.1 Experiment 1: Direction of Overlap

Method

Stimuli: The stimuli consisted of regularly-spaced sawtooth curves as shown in figure 8, which suggested a series of serrated edges overlapping either to the left or the right. The stimulus configurations (figure 9) consisted of ten configurations: three straight-line approximations to concave cusps (shapes 1, 2, and 3), three concave cusps defined by continous curves (shapes 4, 5, and 6), three smoothly convex corners (shapes 7, 8, and 9) and a symmetrical sawtooth (shape 10). Note that the basic sawtooth curve consists of either (i) a cusp feature on the left and a sharp corner on the right, or (ii) a sharp corner on the left and a smoothly rounded corner on the other. The stimuli were displayed as white lines of intensity 5.5 ft.-L. on a grey background of 1.2 ft.-L., in a darkened room. The stimuli were generated by a Symbolics 3670 Lisp Machine and displayed on a Tektronix 690SR color monitor.

The various sawtooth curves were composed of bitmaps. When viewed from 76 inches, one pixel subtended 1'. The .4 mm dot pitch of the CRT permitted individual pixels to be resolved. The smallest continuous cusp (figure 9, shape 4) blended into the straight segment of the sawtooth over an extent of 4 pixels (4'), and the overall amplitude of the sawtooth was roughly 20'. Note that the smallest cusp differed from a sharp corner by the addition of a single pixel (subtending 1'). The difference was just visible from the subject's viewing distance. All configurations (except shape 10 in figure 9) suggest overlap to the right. When a sawtooth stimulus was presented on the display, either the bitmap or its mirror reflection (about the vertical) was projected in order to balance for direction of overlap.

Procedure: The subject's task was to report, by pressing buttons on a mouse (an interactive x-y pointing device), which direction the surfaces overlapped. The task started with a one-second presentation of a fixation point, a cross subtending 8x8'. The fixation point appeared against a blank background at the location where a sawtooth curve would momentarily appear.

AD-A168 613 VISUAL REPRESENTATIONS SUBSERVING TEXTURE PERCEPTION (U) OREGON UNIV EUGENE J BECK ET AL. 15 JAN 86 AFOSR-TR-86-0290 F49620-83-C-0093 2/2

UNCLASSIFIED F/G 5/10 NL

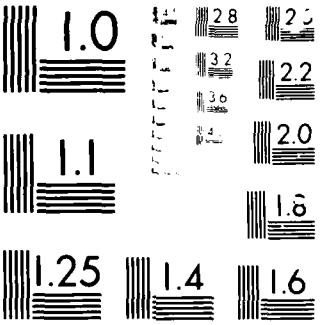
1.0
1.1
1.25
1.4
1.6
2.0
2.2
1.8

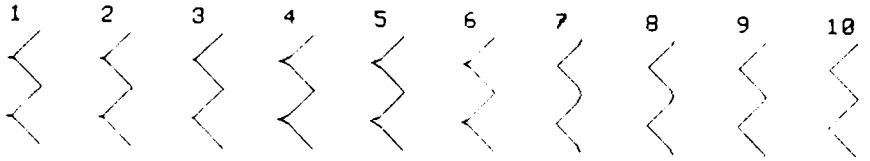

Figure 9. Cusp stimuli for the experiments. Shapes 1-3 are straight-line approximations to cusps. Shapes 4-6 are the continuously curved cusps. Shapes 7-9 are the smooth corners and shape 10 the symmetrical sawtooth.

---

fixation point was placed along a given sawtooth curve at a point midway between two extrema. The sawtooth pattern was then presented for 500 msec during a training run and subsequently at 250 msec during runs for which data were collected. After the stimulus interval or the subject's response, which ever occurred first, the stimulus was replaced by a random dot masking pattern of moderate density chosen to roughly match the mean luminance of the stimulus display. The subject was instructed to make rapid responses of overlap direction while maintaining accuracy. Overlap direction was explained with reference to the edges of the pages of an open book, e.g. the pages on the left "overlap to the left" and if the subject sees this direction of overlap in the stimulus the left mouse button should be depressed. The response time was measured relative to stimulus onset. The experiment consisted a sequence of 100 trials for which overlap direction and reaction times were recorded. The sequence consisted of five repetitions of randomized presentations of the ten stimuli and their mirror reflections. Prior to collecting the data, subjects were given a practice run of 20 trials. The subjects consisted of two females and five males with normal or corrected vision; all were unpaid volunteer subjects.

Results

Table 1 shows the mean and individual error counts for each stimulus shape across subjects. Data was collapsed across mirror reflections, and the error counts for each of the 10 stimulus shapes were computed. A judgment was regarded as an "error" if the subjective direction of overlap did not match that predicted by the shape. Note that for shape 10, the symmetrical sawtooth, there is no correct or incorrect judgment, but that data was included to reveal any subject bias towards interpreting overlap to either the left or right. The mean error would have been 5.0 if no bias existed; the observed mean of 6.2 across subjects shows a bias towards judging overlap towards the right.

Table 1. Error counts by subject.						
Shape	A	B	C	D	E	Mean
1	0	1	1	0	4	1.2
2	0	2	2	0	4	1.6
3	1	3	4	0	4	2.4
4	0	1	0	0	3	0.8
5	0	1	3	0	2	1.2
6	0	0	0	0	5	1.0
7	0	1	2	0	1	0.8
8	0	2	4	0	0	1.2
9	4	2	3	3	4	3.2
10	7	6	6	7	5	6.2

Comparing the results across shape type, the least mean error rate occurred with the continuously-curved cusps (shapes 4-6) in general. Comparing mean error rates within shape or across shapes few differences reached significance. All were significantly different from shape 10 (at the .05 level). Within shape type there was a weakly significant trend for the error rates tended to increase with decreasing size of the feature. Regarding the bias, note that subjects A and D had the strongest bias in shape 10 but the fewest errors, and subject E, with the most errors, had no bias in shape 10. The mean reaction times revealed no significant differences.

Most subjects were able make the judgments both quickly and accurately with little practice. With practice, we found that very accurate judgements were possible with very short presentation times. For example, one of the authors (AB) produced the results shown in table 2 with 80 msec presentation time. Note that shape 10 was not incorporated in this experiment.

Table 2: Subject AB (80 msec)									
Shape	1	2	3	4	5	6	7	8	9
Error counts	0	1	0	0	0	0	0	0	1
Reaction times	428	558	589	438	457	475	526	481	722

Discussion

The first point to draw from this experiment is that all shapes other than the simple corner (shape 10) were roughly equivalent in defining figure-ground. That is, figure-ground can be determined on the basis of convexity in the absence of a cusp (shapes 7-9), or independently, on the basis of a cusp in the absence of convexity (shapes 1-3). Moreover, the straight-line

approximations to concave cusps (shapes 1-3) were similar to the continuously-curved concave cusps (shapes 4-6) despite their visually obvious differences. With regard to the size of the cusp feature for a given shape, the error means varied little between the smallest to largest (e.g. between shapes 1 and 3 or between 4 and 6). The smallest feature was shape 3, which differed from a sharp corner by the addition of a single pixel, was quite effective in defining figure-ground despite it being barely discriminable from a sharp corner (shape 10) at the subject's viewing distance.

The results for subject E are significantly different from those of the other subjects. The task of seeing the contours as edges of overlapping surfaces was difficult for some subjects and required a greater number of learning trials. A few candidate subjects reported that the patterns merely looked like flat, two-dimensional arrangements of lines whose corners and angles, like arrowheads, pointed to either the left or right. For these subjects the impression of the sawtooth curve as a serrated edge was not natural. Consequently we decided to perform an experiment in which apparent figure-ground was probed not by direction of overlap, but more directly by asking whether a dot presented on one or the other side of an indicated curve was on the edge (in the figure suggested by the curve) or to the side, that is, on the background adjacent to the edge. This second task still required subjects to hold an interpretation of the sawtooth curve as the edge of a serrated surface. The subjects' instruction, however, encouraged a more local judgment of the figure-ground relationship of the probe dot relative to the immediately adjacent curve. In reflection on the first experiment, we suspect that although the judgement could be made locally on the basis of whether the indicated curve was a left or right edge, the instructions encouraged a more diffuse, global judgement of whether the pattern represented a cascading series of overlapping edges, and then to report the direction of overlap, a task that some found difficult. The following experiment, it turns out, produced much more uniform results and required fewer introductory trials.

4.4.2 Experiment 2: Judging Figure-ground

Method

Stimuli: The stimuli consisted of patterns of sawtooth curves as in Experiment 1. In addition, each stimulus was presented with an additional dot placed on one or the other side of the sawtooth curve indicated by the fixation point. The dot would be judged as either on or off the given edge, as an independent means of probing apparent figure-ground. The symmetric sawtooth (shape 10) was eliminated in this experiment.

Procedure: Subjects were instructed, as before, to direct attention to the fixation point, then to the relationship between the indicated sawtooth curve and the dot (which appeared randomly to the left or right of the curve). As rapidly as possible while preserving accuracy the subject was then to indicate whether the dot appeared to be on the given edge or not, by depressing a mouse button. The sequence consisted of three repetitions of randomized presentations of the nine stimuli with each of the permutations of overlap and dot position. Prior to data collection,

subjects were given a practice run of one sequence of 36 trials; data was then collected for a sequence of 108 trials (3 sequences of 36). Again the viewing distance in all cases was 76 inches, so that one pixel subtended 1′. Four female and five male graduate students, all unpaid volunteers, participated as subjects.

Results

Table 3 shows the mean and individual error counts for each stimulus shape across subjects. Again data was collapsed across mirror reflections, and the error counts for each stimulus shape was computed.

Table 3. Error counts by subject

Shape	A	B	C	D	E	F	G	H	I	Mean
1	0	0	1	1	2	1	1	1	2	1.0
2	1	0	1	1	0	2	2	0	2	1.00
3	0	0	0	4	1	2	1	0	0	0.89
4	0	0	0	0	1	1	1	2	0	0.56
5	0	0	0	1	0	1	1	1	0	0.44
6	1	0	2	2	1	1	0	0	1	0.89
7	0	1	0	1	0	4	0	0	0	0.67
8	0	0	0	0	0	2	1	0	0	0.33
9	0	2	3	6	0	2	1	2	2	2.00

Pooling across type (i.e. shapes 1-3, 4-6, and 7-9) revealed no significant differences. Within type, the straight-line approximated cusps (shapes 1-3) were insignificantly different, as were the continuously-curved cusps (shapes 4-6). We found the convex shape 9 (having the smallest radius of curvature) to be different from the other convex shapes (shape 8, at the .05 level and shape 7 between the .1 and .05 level).

Discussion

Generally one is struck more by the similarity of the results across type and within type than with the differences. Also, in light of the very low error rates for all conditions tested (where the worst mean of 2 out of 12 occurred for shape 9) shows that the visual system very efficiently uses minute evidence along a curve in determining figure-ground. The only differences concerned shape 9, the smallest of the rounded corners. Seemingly, the introduction of even a slight amount of curvature to the sawtooth pattern biases apparent figure-ground. With reference to figure 9, note that the curvature is induced by only 2 to 3 pixels of "rounding" which subtends only 2-3′. Likewise, the introduction of a minute cusp-like feature is effective at nearly the limit of visual resolution, such as the single pixel in shape 3, which amounts to a 1′ tip added to convert the sharp corner to a minute crack-like cusp subtends a similarly small locality.

4.5 General Discussion

Starting from earlier observations that convexity, as a general property, influences apparent figure-ground, we have examined how convexity might be measured. Earlier we pointed out (section 2.3) that convexity has many geometric definitions in principle, each suggesting visual algorithms of varying tractability and utility. We wish to emphasize here that global measures of convexity, particularly those that require an isolated and closed contour are probably inappropriate, and that instead local evidence for figure-ground (still based on convexity) would be preferable.

We proposed that the concave cusp discontinuity would be a useful feature of a curve on which to base an early decision of figure-ground. The experiments we performed confirmed the proposal, to the extent that our stimuli successfully isolate the concave cusp from the already-expected contribution of curvature convexity. We designed shapes 4-6 as representative examples of concave cusps, shapes 7-9 as representative examples of convexity without cusps, and shapes 1-3 as examples of cusps without convexity. We designed a straight-line approximation to a cusp, to see if the particular local geometry, and not the curvature of individual curved arcs, may induce the figure-ground decision. We found that all three shape types were roughly equally effective, from which we conclude that while convexity alone is sufficient (shapes 7-9), as expected, so is the distinctive geometry of the concave cusp, even when one contrives to minimize the contribution due to contour curvature (shapes 1-3 as opposed to shapes 4-6). Recall the arguments to this effect in the demonstrations involving figures 7.

The concave cusp seemingly behaves as a geometrically-defined feature, and while we do not offer a concise definition, the apparent equivalence of the cusps composed of straight-line segments and those composed of smooth arcs suggests the definition is rather primitive — the component arcs need not curve continuously in the vicinity of the cusp for the arrangement to be effective (refer back to figures 7b and 8c as well). Earlier we noted that the concave (negative-negative) cusp is but one of six arrangements of arcs at a discontinuity of tangent along a curve, and predicted that it should be the most effective. Variations such as those suggested in figure 4 should permit a more specific description of what geometric aspects constitute this figure-ground determiner.

REFERENCES

Andrews, D.P. 1967a Perception of contour orientation in the central fovea. Part I. Short lines. *Vision Research* **7**, 975-997.

Andrews, D.P. 1967b Perception of contour orientation in the central fovea. Part II. Spatial integration. *Vision Research* **7**, 999-1013.

Attneave, F. 1955 Perception of place in a circular field. *American Journal of Psychology* **68**, 69-82.

Attneave, F. 1974 Apparent movement and the what-where connection. *Psychologia* **17**, 108-120.

Attneave, F. 1982 Praegnanz and soap bubble systems : a theoretical exploration. In *Organization and Representation in perception.*, Beck, J., ed. 11-29.

Bacon, J. and King-Smith, P.E. 1977 The detection of line segments. *Perception* **6**, 125-131.

Beck, J. & Halloran, T. 1985 Effects of spatial separation and retinal eccentricity on two-dot vernier acuity. *Vision Research* **25**, 1105-1111.

Beck, J. Effect of orientation and of shape similarity on grouping. *Perception and Psychophysics* **1**, 300-302, 1966.

Beck, J. Similarity grouping and peripheral discriminability under uncertainty. *American Journal of Psychology* **85**, 1-19, 1972.

Beck, J. Textural segmentation. In Beck, J. (Ed.) *Organization and representation in perception.* Hilsdale, NJ: Lawrence Erlbaum Associates, 1982, 285-317.

Beck, J. Textural segmentation, second-order statistics, and textural elements. *Biological Cybernetics* **48**, 125-130, 1983.

Beck, J., Prazdny, K. & Rosenfeld, A. A theory of texture segmentation. In Beck, J., Hope, B. & Rosenfeld, A. (Eds.) *Human and machine vision.* New York: Academic Press, 1983, 1-38.

Burton, G.J. & Ruddock, K.H. 1978 Visual adaptation to patterns containing two-dimensional spatial structure. *Vision Research* **18**, 93-99.

Caelli, T.M. & Julesz, B. 1978 On perceptual analyzers underlying visual texture discrimination: Part I. *Biological Cybernetics* **28**, 167-175.

Caelli, T.M., Preston, G.A.N., & Howell, E.R. 1978 Implications of spatial summation models for processes of contour perception: a geometric perspective. *Vision Research* **18**, 723-731.

De Valois, R.L., Albrecht, D.G., & Thorell, L.G. 1982 *Vision Research* **22**, 545-559.

Enroth-Cugell, C. & Robson, J.G. 1966 The contrast sensitivity of retinal ganglion cells of the cat. *J. Physiol.* **187**, 517-552.

Frisch, H.L, & Julesz, B. 1966 Figure-ground perception and random geometry. *Perception and Psychophysics* **1**, 389-398.

Glass, L. 1969 Moire effect from random dots. *Nature* **223**, 578-580.

Glass, L. 1979 Physiological mechanisms for the perception of random dot moire patterns. In *Pattern formation by dynamic systems and pattern recognition.* H. Haken, ed. Berlin: Springer-Verlag.

Glass, L. & Switkes, E. 1976 Pattern recognition in humans: correlations which cannot be perceived. *Perception* **5**, 67-72.

Hochberg, J. 1964 *Perception.* New Jersey: Prentice-Hall.

Hoffman, D.D. 1983 The interpretation of visual illusions. *Scientific American* **249**, December.

Hoffman, D.D. & Richards, W. 1982 Representing smooth plane curves for recognition: Implementations for figure-ground reversal. *Proceedings of the National Conference of the American Association for Artificial Intelligence*, 5-8.

Hoffman, D.D. & Richards, W. 1984 Parts of recognition. *Cognition* **18**, 65-96.

Howell, E.R. & Hess, R.F. 1978 The functional area for summation to threshold for sinusoidal gratings. *Vision Research* **18** 369-374.

Hubel, D.H. & Wiesel, T.N. 1962 Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106-154.

Hubel, D.H. & Wiesel, T.N. 1968 Receptive fields and functional architecture of monkey striate cortex. *J. Physiol.* **195**, 215-243.

Janez, L. 1984 Visual grouping without low spatial frequencies. *Vision Research* **24**, 271-274.

Julesz, B. 1965 Texture and visual perception. *Scientific American* **212**, 38-48.

Julesz, B. Perceptual limits of texture discrimination and their implication for figure-ground perception. In Leewenberg, E. & Buffart, H. (Eds.) *Formal theories of perception.* New York: Wiley, 1978, 205-216.

Julesz, B. 1981 Figure and ground perception in briefly presented isodipole textures. In *Perceptual organization*, Kubovy, M. & Pomerantz, J.R., eds. Hillsdale, New Yersey: Lawrence Erlbaum Associates, 27-54.

Julesz, B. & Schumer, R. Early visual perception. *Annual Review of Psychology* **575-627**, 1981.

Kaniza, G. 1979 *Organization in Vision.* New York: Praeger.

Kaniza, G. & Gerbino, W. 1976 Convexity and symmetry in figure-ground organization. In *Art and Artefacts*, M. Henle, ed. 25-32. New York: Springer.

Koehler, W. 1929 *Gestalt Psychology.* New York: H. Liveright.

Koffka, K. 1935 *Principles of Gestalt Psychology.* New York: Harcourt, Brace.

Marr, D. 1976 Early processing of visual information. *Phil. Trans. R. Soc. Lond. B* **275**, 483-524.

Marr, D. 1982 *Vision: a computational investigation into the human representation and processing of visual information.* San Francisco: Freeman.

Marr, D. & Hildreth, E. 1980 Theory of edge detection. *Proc. Roy. Soc. Lond. B.* **207**, 187-217.

Marr, D. & Nishihara, H.K. 1978 Representation and recognition of the spatial organization of three-dimensional shapes. *Proc. R. Soc. Lond.* **211**, 151-180.

Marr, D., Poggio, T. & Hildreth, E. 1980 Smallest channel in early human vision. *J. Opt. Soc. Am.* **70**, 868-870.

Marroquin, J.L. 1976 Human visual perception of structure. M.S. Thesis. Department of Electrical Engineering and Computer Science, MIT.

Movshon, J.A., Thompson, I.D. & Tolhurst, D.J. 1978 Receptive field organization of complex cells in the cat's striate cortex. *J. Physiol., Lond.* **283**, 79-99.

O'Callaghan, J.F. 1974a Computing the perceptual boundaries of dot patterns. *Computer Graphics and Image Processing* **3**, 141-162.

O'Callaghan, J.F. 1974b Human perception of homogeneous dot patterns. *Perception* **3**, 33-45.

O'Callaghan, J.F. An alternative definition for "neighborhood of a point". *IEEE Trans. Computers*, Nov., 1121-1125.

Poggio, G.F. 1972 Spatial properties of neurones in striate cortex of unanesthetized macaque monkey. *Investve Ophth.* **11** 368-377.

Prazdny, K. 1984 On the perception of Glass patterns. *Perception* **13**, 469-478.

Richards, W. & Purks, S.R. 1978 On random-dot texture discrimination. *J. Opt. Soc. AM.* **68**, 270-271.

Rubin, E. 1958 Figure and Ground. In *Readings in perception*, D.C. Beardslee and M. Wertheimer, eds. Princeton, N.J.: Van Nostrand.

Schiller, P.H., Findlay, B.L. & Volman, S.F. Quantitative studies of single-cell properties in

monkey striate cortex. I. Spatiotemporal organization of receptive fields. *J. Neurophysiol.* **39**, 1288-1319.

Scobey, R.P. 1982 Human visual orientation discrimination. *Journal of Neurophysiology* **48**, 18-26.

Stevens, K.A. 1978 Computation of locally parallel structure. *Biological Cybernetics*

Stevens, K.A. 1986 The concave cusp as a determiner of figure-ground. In preparation.

Treisman, A. Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance* **8**, 194-214, 1982.

Ullman, Shimon 1979 *The interpretation of visual motion.* Cambridge, MA: MIT Press.

Ullman, S. 1984 Visual Routines. *Cognition* **18**, 97-159.

Vassilev, A. & Penchev, A. 1976 Spatial and temporal summation in the perception of lines. *Vision Research* **16**, 1329-1335.

Waltz, D.I. 1975 Generating semantic descriptions from drawings of scenes with shadows. Ph.D. dissertation, AI Lab, MIT, 1972.

Watt, R.J. & Morgan, M.J. 1983 The recognition and representation of edge blur: evidence for spatial primitives in human vision. *Vision Research* **13** 1465-1477.

Wertheimer, M. 1923 Untersuchungen zur Lehre von der Gestalt. *Psych. Forch.* **4**, 301-350. Abridged translation: Principles of perceptual organization. In *Readings in Perception*, 1958, D.C. Beardslee & M. Wertheimer, Eds. New York: D. Van Nostrand.

Wilson, H.R. & Bergen, J.R. 1979 A four mechanism model for spatial vision. *Vision Research* **19**, 19-32.

Wright, M.J. 1982 Contrast sensitivity and adaptation as a function of grating length. *Vision Research* **22**, 139-149.

Zucker, S.W. 1983 Computational and psychophysical experiments in grouping: early orientation selection. In *Human and machine vision*, Beck, J., Hope, B., and Rosenfeld, A., eds. 545-567.

PUBLICATIONS, MEETINGS AND PERSONNEL

Publications:

Beck, J., Prazdny, K. & Ivry, R. 1984 The perception of transparency with achromatic colors. *Perception and Psychophysics* **35**, 407-422.

Beck, J. 1985 Perception of transparency in man and machine. *Computer Vision, Graphics and image Processing* **31**, 127-138.

Beck, J. & Halloran, T. 1985 Effects of spatial separation and retinal eccentricity on two-dot vernier acuity. *Vision Research* **25**, 1105-1111.

Beck, J., Sutter, A. & Ivry, R. 1986 Spatial filtering, features, and grouping in texture segmentation. *Computer Vision, graphics and image processing.* (In preparation).

Beck, J. & Ivry, R. 1986 Aggregation of texture elements on the basis of geometric descriptors. *Perception and Psychophysics.* (In preparation).

Stevens, K.A. 1984 On gradients and texture "gradients". Commentary on: Cutting & Millard 1984. Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General* **113**, 217-220.

Stevens, K.A. & Brookes, A. 1986 The convex cusp as a determiner of figure-ground. Submitted to *Perception*

Stevens, K.A. 1985 Effect of defocussing on the detection of collinear structure. In preparation.

Stevens, K.A. & Lulich, D.P. 1986 Artifacts in the perception of fine detail. In preparation.

Meetings:

Beck, J. "Two-dot vernier acuity", Psychonomic Society. San Diego, November, 1983.

Beck, J. "Vernier acuity and virtual lines", Workshop on Visual Structure, Fairchild Camera and Instrument corporation. Palo Alto. February, 1984

Beck, J. "Texture segmentation", Workshop on Perceptual Organization, Pajaro Dunes, June, 1984

Beck, J. "Perception of transparency in man and machine", Second Human and Machine Vision Workshop, Montreal, Canada, August, 1984

Beck, J. "Sensory, organizational, and inferential processes in lightness perception", APA, Toronto, Canada, August, 1984

Beck, J. "Grouping processes in texture segmentation", Psychonomic Society. San Antonio. November, 1984

Beck, J. "Grouping and texture segmentation", Center for Adaptive Systems, Boston University, April, 1985; Department of Psychology, University of Maryland, April, 1985; CAS Research, Schlumberger, Palo Alto, May, 1985; University College (London), July, 1985.

Stevens, K.A. "Shape predicates and texture gradient psychophysics". Invited talk, Vision Review. Endicott House, Massachusetts Institute of Technology. January, 1984.

Stevens, K.A. "The case for place tokens" Workshop on Visual Structure. Fairchild Camera and Instrument Corporation. February, 1984.

Stevens, K.A. "Early representations of visual texture: symbolic versus analogic". Invited talk, Department of Computer Science, University of British Columbia. April, 1984.

Stevens, K.A. "Visual primitives and their aggregation processes in texture". Review of Sponsored Research. Air Force Office of Scientific Research. Sarasota, Florida. May, 1984.

Stevens, K.A. "Semi-invariants over back-projection". Invited talk, Workshop on Perceptual Organization. Pajaro Dunes. June, 1984.

Stevens, K.A. "Inferring gaussian curvature from curves across surfaces". Optical Society of America. San Diego, California. October, 1984.

Professional personnel associated with the research:

Faculty:

Prof. Jacob Beck
Prof. Kent A. Stevens

Graduate students:

Allen Brookes
William Goodwin
Richard Ivry
Cathryn Stanford
Anne Sutter

Undergraduates:

Kurt Benson
Chrisopher Manzer
Richard Wildes

END

DTIC

7-86